On Limits on the Provable Consequences of Quantum Pseudorandomness

Samuel Bouaziz--Ermann*[1], Minki Hhan†[2], Garazi Muguruza[3], Quoc-Huy Vu[4]

[1]Okinawa Institute of Science and Technology Graduate University, Japan
samuel.bouazizermann@oist.jp
[2]KAIST, Daejeon, Korea
minkihhan@kaist.ac.kr
[3]Surf, The Netherlands
garazimugu@gmail.com
[4]De Vinci Higher Education, De Vinci Research Center, Paris, France
quoc.huy.vu@ens.fr

Abstract

There are various notions of quantum pseudorandomness, such as pseudorandom unitaries (PRUs), pseudorandom state generators (PRSGs), pseudorandom function-like state generators (PRFSGs), and quantum-computable pseudorandom generators (QPRGs). Unlike the different notions of classical pseudorandomness, which are known to be existentially equivalent to each other, the relations among quantum pseudorandomness have yet to be fully established.

We present evidence suggesting that some forms of quantum pseudorandomness are unlikely to be constructed from the others. This indicates that quantum pseudorandomness behaves quite differently from classical pseudorandomness.

Our main result is a unitary oracle separation where log-length output PRFSGs exist but QPRGs with negligible correctness error do not. This result suggests that the inverse-polynomial error in the state-of-the-art construction of QPRGs from log-length PRSGs is inherent. To achieve this, we prove a novel differential geometric *barrier theorem* for the product Haar measure on quantum states which replaces the usual concentration inequalities by certifying a non-negligible "gap" between two large separated sets.

Our separation is based on an oracle that outputs Haar random quantum states for each bit string, which can be viewed as a quantum state version of the random oracle model, where output strings are replaced by quantum states.

Variations of this oracle can be used to study other relationships between quantum cryptographic primitives, and we use it to achieve partial separations, which highlight technical difficulties when dealing with ancillary registers, measurements, and adaptivity in the quantum setting. More precisely, we obtain separations showing limitations of: (i) deriving ancilla/measurement-free pseudorandom unitaries (PRUs) from PRFSGs, and (ii) a natural way of constructing super-log-length output PRSGs from log-length output PRFSGs. The latter partly complements the known hardness of shrinking the PRSG output lengths.

*This work was done while the author was affiliated at Sorbonne Université, CNRS, LIP6, France.
†This work was done in part while the author was affiliated with KIAS, Korea, and UT Austin, USA.

Contents

1 Introduction **3**
- 1.1 Our results . . . 4
- 1.2 Related works . . . 6

2 Technical overview **8**
- 2.1 Separating QPRGs from short PRFSGs . . . 8
- 2.2 On separating PRUs from PRFSGs . . . 10
- 2.3 On the length extension of PRSGs . . . 11

3 Preliminaries **11**
- 3.1 Quantum states, operations, and trace . . . 11
- 3.2 Haar random states and unitaries . . . 12
- 3.3 Cryptographic primitives . . . 13
- 3.4 QPSPACE oracle . . . 14
- 3.5 State property tests . . . 14
- 3.6 Quantum OR lemma . . . 17
- 3.7 Useful lemmas . . . 17

4 Common Haar Function-like State Oracles **18**
- 4.1 CHFS oracles and unitarization . . . 18
- 4.2 Construction of PRFSGs in the CHFS model . . . 19

5 Oracle Separation of QPRGs from PRFSGs **21**
- 5.1 The barrier theorem . . . 22
- 5.2 Impossibility of QPRGs . . . 25
- 5.3 Proof of Lemma 5.4 . . . 27

6 On Separating PRUs from PRFSGs **34**

7 On Length Extension of PRSGs **38**

A Formal Proof of the Barrier Theorem **44**
- A.1 Removing differentiability assumption . . . 44

B Average-case Hard Problem **44**

1 Introduction

In classical cryptography, computational pseudorandomness generated by pseudorandom generators (PRGs) and functions (PRFs) serves as a central resource. It can be used for many applications, such as commitments [Nao91], digital signatures [Rom90], and symmetric key encryptions. Furthermore, the existence of PRGs and PRFs is necessary for the existence of almost all cryptographic primitives with computational security, including one-way functions (OWFs) [HILL99].

However, when we treat the world as operating under the laws of quantum mechanics, the notion of pseudorandomness must be revisited. Ji, Liu, and Song [JLS18] proposed the first two inherently quantum pseudorandom primitives, pseudorandom state generators (PRSGs) and unitaries (PRUs). Quantum pseudorandomness has been shown to be useful for constructing many (quantum) cryptographic primitives, for example, PRSGs imply quantum commitments and oblivious transfers [MY22, AQY22]. On the other hand, Kretschmer showed that PRSGs and PRUs (with super-logarithmic output lengths) are potentially weaker primitives than classical pseudorandomness by presenting an oracle separation [Kre21]. The dramatic interest in fundamentally quantum cryptographic primitives emerged as a new cryptographic direction possible even in the world without one-way functions.

Quantum pseudorandomness turns out to be quite different from its classical counterparts. To begin with, the length becomes an important parameter in quantum cryptography. Given an λ-bit input seed, [BS20] shows that $c\log\lambda$-length output PRSGs exist unconditionally for $c \ll 1$ whereas achieving $c \geq 1$ requires computational assumptions. We name $c\log\lambda$-length and superlogarithmic-length PRSGs by short PRSGs and long PRSGs, respectively. [BM24, CGG24] shows that long PRSGs cannot be used to construct short PRSGs, highlighting that the output lengths of PRSGs cannot be shrunk. Note that in the classical setting, the output length of PRGs can easily be shrunk and lengthened arbitrarily.

Contrary to the long PRSGs that are separated from classical cryptography, [ALY24] shows that short PRSGs can be used to construct a variant of PRGs called pseudodeterministic quantum-computable PRGs (QPRGs) with an inverse-polynomial pseudodeterminism error. A QPRG G with an ε error is a QPT algorithm that, on $(1-\varepsilon)$-fraction of the input seed, outputs the same value with probability $1-\varepsilon$.[1] When ε is negligible, we simply call G a QPRG. With an inverse-polynomial ε, one can construct digital signatures and IND-CPA encryptions [BBO+25], although the proof technique is more complicated than the classical one.

[ALY24] questioned if the error ε of QPRG can be reduced to be negligible. If it is possible, virtually classical Minicrypt can be recovered using the same technique. Thus, our main question is:

Do short PRSGs imply QPRGs (with negligible errors)?

In particular, if the answer is affirmative, long PRSGs can be constructed from short PRSGs through QPRGs by using known constructions [GGM86, JLS18]. Yet, if the answer is no, there is still a possibility of lengthening PRSG outputs without relying on classical cryptographic primitives, thus we ask:

Can we extend the length of PRSGs?

We also consider another difference between classical and quantum in the landscape of pseudorandomness. For example, we do not know how to construct PRUs from PRSGs or even pseudorandom function-like state generators (PRFSGs) [AQY22, AGQY22], whereas classically PRGs can be used to construct PRFs and vice versa. We are left with an unsatisfactory state of affairs; unlike in classical pseudorandomness, there is no single assumption unifying quantum pseudorandomness. This makes us ask:

Are PRSGs, PRFSGs, and PRUs existentially equivalent?

[1] The second error (for the same seed) can be arbitrarily reduced by the repetition. We use this notion following the original definition [ALY24].

1.1 Our results

We provide negative evidence for the above questions by presenting new oracle worlds where one primitive exists, but the other does not exist, or at least is hard to construct, by showing that some natural constructions are insecure. Our results suggest that quantum pseudorandomness could behave fairly differently from classical pseudorandomness.

Common Haar function-like state model. Our separations are based on variants of the common Haar function-like state (CHFS) oracles where for each input $x \in \{0,1\}^*$ the oracle outputs a Haar[2] random state $|\phi_x\rangle$ of length $\ell(|x|)$ where $|x|$ is the bit-length of x. Note that this is an isometry. We also consider the unitary variants that instantiate this oracle.

Since the construction of PR(F)SGs with output length $\ell(|x|)$ is straightforward with these oracles, our main contribution is to show that many other primitives are hard to construct even with the CHFS oracle (or its variants).

Separating QPRGs from short PRFSGs. Using the CHFS oracles, our main contribution is to give a negative answer to the open question asked in [ALY24].[3] More formally, we prove the following theorem.

Theorem 1.1. *For $c\log\lambda \le L \le C\lambda$ for some $c, C > 0$, there exists a unitary oracle relative to which PRFSGs with L-qubit outputs exist but QPRGs do not.*

Note that with the negligible errors, the classical constructions work well even for the quantum-computable counterparts. Therefore, this theorem states that QPRGs, as well as quantum-computable one-way functions (QOWFs) and quantum-computable pseudorandom functions (QPRFs), constructed from quantum primitives must suffer from inverse polynomial errors, as in the typical constructions that uses tomography.

This theorem has several interesting implications. First, it complements the classical–quantum cryptography separation of [Kre21] along two incomparable dimensions. On the classical side, [Kre21] considers primitives with perfect correctness (in particular, OWFs)[4], whereas our separation continues to hold even when the primitive is allowed to have a negligible correctness error. On the quantum side, [Kre21]'s separation requires long output length (in particular, $\omega(\log\lambda)$-length PRUs), while our oracle separation works even for short PRSGs. Thus, our result shows that the gap between classical and quantum cryptography persists even when correctness is relaxed and the quantum pseudorandomness assumption has a significantly shorter output length.

Secondly, it can be interpreted as finding a quantum-computable version of Pessiland, where there are hard-on-average languages in $\mathbf{QCMA} \cap \mathbf{coQCMA}$ yet there are no QPRGs. Previously, [BNY25] gave a black-box separation of PRGs from EV-OWPuzz, thereby separating PRGs and hard-on-average languages in **QCMA** (from EV-OWPuzz).[5] This result strengthens the result to include **coQCMA** and to unitary oracle separation. To show this, we choose $\ell = O(\log\lambda)$ and consider a state version of the decisional permutation inversion problem [BBBV97] that decides if the given input (str, z) satisfies $x \le z$ where str specifies a state $|\psi_{\mathsf{str}}\rangle$ that is sufficiently close to the CHFS oracle's output state $|\phi_x\rangle$ for some x. By appropriately choosing the output length ℓ, the uniqueness of x is ensured with overwhelmingly high probability. We outline this language and interpretation in Appendix B. In the classical setting, [Wee06] found a Pessiland oracle.[6]

Corollary 1.2. *There exists a unitary oracle relative to which a quantumly hard-on-average promise problem in $\mathbf{QCMA} \cap \mathbf{coQCMA}$ exists, but no quantum-computable one-way functions or pseudorandom generators exist.*

[2] We sometimes choose the oracles from its sub-distribution, but we stick to use the name CHFS oracles for simplicity.

[3] [Bar26] concurrently resolves this question by showing a black-box impossibility. See the concurrent work for more detailed discussion.

[4] Technically, they showed the stronger statement $\mathbf{BQP} = \mathbf{QMA}$.

[5] Similarly, the worst-case hardness $\mathbf{BQP} \neq \mathbf{QCMA}$ was proven in [BM24, CGG24].

[6] The authors noted that an unpublished work by Impagliazzo and Rudich gave the first Pessiland oracle.

On the other hand, however, we found that this oracle world is not too pessimistic. Recall that PRSGs with output length $c \log \lambda$ for large enough c can be used to construct digital signatures and IND-CPA encryptions, with classical keys and outputs, using the notion of recognizable aborts [BBO+25]. This gives the following corollary in the world of quantum-computable classical-communication (QCCC) cryptography [CGG24], which was separated from (classical-computable) classical cryptography in [KQT25].[7] This strengthens a similar black-box separation proven in [BNY25].

Corollary 1.3. *Relative to the same oracle, quantum-computable digital signatures and IND-CPA symmetric-key encryptions exist, but no quantum-computable OWFs or PRGs exist.*

We found this interesting, as it shows that OWFs and PRGs are *stronger* than the main applications, encryptions and signatures. This is possible because internal randomness in quantum machines cannot be extracted as random coins, unlike randomness in classical machines as observed in several previous works [AA11, AIK22, BCQ23, BNY25]. Consequently, our separation also separates EV-OWPuzz from QOWFs, confirming the belief of [CGG24] that EV-OWPuzz could be the central primitive for QCCC cryptography.[8] We refer the reader to Fig. 1 for further implications that can be derived from our main theorem.

Further results: Partial separations. Variations of the oracle used to prove Theorem 1.1 can be used to study the relationship of other quantum pseudorandom primitives, for which a complete separation or construction remains out of reach, and appear to require additional ideas beyond the current techniques. We include these partial separations to indicate where the obstructions come from, as this show what is left to prove for a full separation, or on the other hand, what a construction needs to overcome.

On constructing PRUs from PRFSGs, without ancilla registers. We study the hardness of constructing PRUs from PRFSGs. We prove that some restricted candidate PRU constructions fail to be secure. The formal statement is as follows.

Theorem 1.4 (Restricted PRU separation)**.** *There exists a unitary oracle*[9] *relative to which adaptively-secure quantum-accessible PRFSGs exist, but non-adaptively secure (and inverseless) PRUs whose implementation uses neither ancillary registers nor intermediate measurements do not.*

Our oracle consists of the CHFS oracle for $\ell(|x|) = |x|$ and the **QPSPACE** oracle that computes the unitary polynomial-space circuit given as input. Our impossibility shows that even the strongest form of PRFSGs cannot be used to construct the weakest form of PRUs in a black-box way without using ancillary registers and measurements.[10] We believe the black-box separation between PRFSGs and PRUs *without the no-ancilla and measurement conditions* also holds in the same oracle world, and leave this open question for future work.

On the length extension of PRSGs. We now turn to the problem of extending the output length of PRSGs. There are already many approaches with some partial positive answers to this problem that we will summarize below. However, we believe that there is no general black-box construction of long-PRSGs from short-PRSGs. We note that proving such a black-box impossibility result would imply Theorem 1.1, as long-PRSGs can be built from QPRGs, and thus studying this separation is a natural next step. We show some partial impossibility result regarding this question, that highlights its difficulty. We consider another natural class of length extensions, including the extension algorithm that takes small PRSs non-adaptively as input and applies a unitary (e.g., $G_k \to U_k(|\phi_1\rangle \otimes |\phi_2\rangle)$ for smaller PRSs $|\phi_1\rangle, |\phi_2\rangle$).

[7] More precisely, [KQT25] constructed the oracle world relative to which $\mathbf{P} = \mathbf{NP}$ yet quantum-computable trapdoor one-way functions exist.

[8] We note that [CGG24] does not exclude the possibility that QOWFs are equivalent to EV-OWPuzz. We refute this possibility at least in a black-box way.

[9] In this paper, we assume that the algorithms can access unitary oracles and its inverses. We do not consider the controls, conjugates or transposes of the oracles, but we believe our results can be extended to them using a similar idea from [Zha25].

[10] The impossibility is shown by an explicit adversary that only uses PRU generation algorithms non-adaptively, so the other forms of PRUs are automatically impossible to construct without ancilla registers and measurements.

Theorem 1.5 (Conditional PRSGs length-extension separation)**.** *There exists an isometry oracle relative to which short PRFSGs exist but long PRSGs whose generation algorithm makes non-adaptive oracle queries and uses no ancilla, do not.*

In fact, our result is slightly stronger: any PRSG length extension of this form is impossible.[11] Moreover, the impossibility of long PRSGs also holds with constant-size ancilla for a certain type of algorithm that revert the ancilla to 0. We note that there are still multiple ways to extend the length of PRSGs that our result does not cover. Besides using superconstant-size ancilla, two notable approaches[12] showing that some PRSG length extensions are *possible* as follows:

- Construct QPRGs first using tomography [ALY24] (with inverse polynomial errors) and then use them to construct new PRSGs. In [BNY25], the authors show that the PRSG length extension is possible in the log-length regime, albeit with quantum key sampling (i.e., the keys are not uniformly distributed but quantumly sampled). This strategy is excluded from our restricted model. Partial progress to construct long PRSGs from short PRSGs was also discussed in the same paper.
- Use quantum queries to the oracle. This is excluded because of our classical-accessible oracle model. In fact, for a length-ℓ PRSG $\{\phi_k\}$, the state $|0\rangle|\phi_a\rangle+|1\rangle|\phi_b\rangle$ for two random keys a, b forms a length-$(\ell+1)$ PRSGs for $\ell=\omega(\log\lambda)$.

We believe the above strategies, allowing length extension up to log-length, are optimal.[13] The full impossibility of PRSG length extension would complement the impossibility of shrinking the output length of PRSGs [BM24, CGG24], and suggests that both primitives are in fact incomparable. Moreover, given the construction of one-way state generators (OWSGs) from short PRSGs [MY24a, CGG+25], it provides evidence for the hardness of constructing PRSGs from OWSGs, while the other direction is possible [CGG+25]. We believe that the PRSG length extension is the most challenging problem discussed in this paper, and any progress on this problem seems to give new interesting techniques.

A summary of our results is given in Fig. 1.

1.2 Related works

Quantum oracle models. The isometry CHFS oracle was first studied in [AGL24] to show the (isometry oracle) separation between QCCC primitives and PRSGs. Recent works suggest different quantum oracle models. The common Haar state model [CCS25, AGL24] represents a world where copies of a random single Haar random quantum state are easily generated. A similar model without restricting Haar randomness was also studied in [DLS24, MY24b, Qia24]. The quantum Haar random unitary oracle model (QHROM) was suggested in [CM24, BFV20], and the applications are studied in [ABGL25, HY25].

Quantum black-box impossibility. Recently, various black-box impossibilities have been shown based on new oracles and new techniques. We briefly summarize this line of research.[14] The separation between OWFs and quantum primitives relative to a quantum oracle [Kre21] initiated this direction, and the same oracle later was shown to imply the hardness of shrinking PRSG output lengths [BM24, CGG24]. This result was later strengthened relative to a classical oracle [KQST23] albeit for weaker quantum primitives. A separation between classical and quantum-computable OWFs is shown in [KQT25].

[11] For example, for $s>t$ and $s=\omega(\log\lambda)$, PRSGs with output length s cannot be constructed from PRSGs with output length t if the longer PRSGs follow the described algorithms.

[12] We remark that there is another recent work [LV24] that discusses the possibility of the length extension of the PRSGs, but only for very specific forms. Furthermore, their work only shows how to do length extension from PRSGs with *super-log* output size.

[13] More explicitly, the first approach gives PRSGs with any log-length output using classical queries to the other log-length output PRSGs. The second approach gives a length-$s+O(\log\lambda)$ output PRSG given quantum access to the other length-s output PRSG.

[14] For the full relations, we refer Microcrypt-zoo.

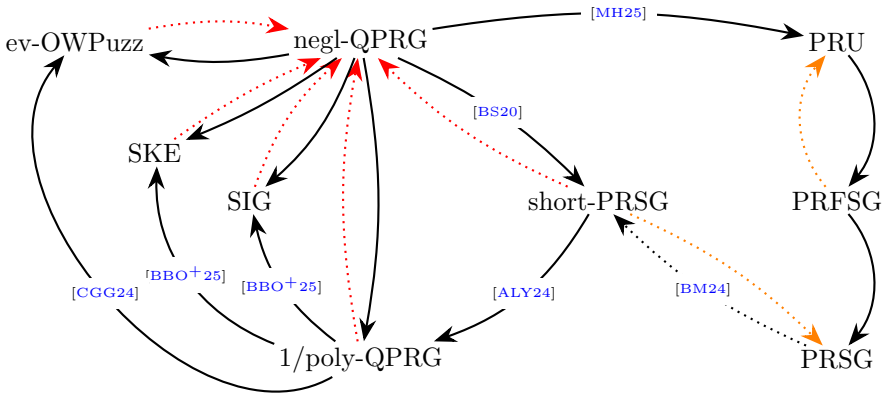

Figure 1: Implications between primitives are represented with an arrow, and separations with a dotted arrow. Our results are the arrows in red and orange, where red arrows indicate an oracle separation and orange arrows indicate conditional/restrictive black-box separations.

Relative to the common Haar state oracles, various separations are implied, e.g., commitments (and EFI pairs [BCQ23]) and single-copy PRSGs exist but no OWSGs and (multi-copy) PRSGs [CCS25, BMM+25, BCN25]. In [BMM+25], they also show a black-box separation between quantum money and EFI pairs.

The isometry version of CHFS oracle provides a world with PRFSGs but without QCCC primitives [AGL24]. On the other hand, an oracle world with QCCC key exchange where **BQP** = **QCMA** holds was introduced in [GMMY26], with some more separations.

Finally, a very recent work [BNY25] shows the black-box impossibility of constructing OWSGs from $\perp$-(Q)PRGs (that can be seen as a weaker version of the QPRGs with negligible correctness errors, see [BBO+25]). A difference between quantum sampling of the keys and uniformly random keys is also explored in the same paper; in this work, we only assume the uniform key setting.

Concurrent work. A concurrent and independent work [GLMY25] shows the oracle separation between PRFSGs and PRUs using similar oracles but with different techniques. The full separations remain open as both papers consider the bounded-length ancilla. We, however, note that [GLMY25] achieves a stronger result, as it separates PRFSG from PRU with a logarithmic number of ancilla qubits, whereas our paper requires no ancilla implementation. They also consider the separations regarding the pseudorandom isometries [AGKL24]. The results about the log-length CHFS oracles are unique to this paper.

[Bar26] concurrently addresses the main question of this paper. They consider a black-box separation between QPRGs and PRSGs. Under their definition, the tomography-based construction of [ALY24] is not treated as black-box, whereas our unitary oracle separation rules out such a tomography approach as well. Ours also exclude the use of PRFSGs. On the other hand, [Bar26] proves the separation regarding $\perp$-PRG and QPRG which we did not.

Acknowledgments. We thank Takashi Yamakawa for the helpful discussion on the definition of **QPSPACE** oracle, and Shoga Yamada and Yao-Ting Lin for noting a minor error in the previous version. MH is supported by Schmidt Sciences Polymath award to David Soloveichik. SBE is supported by PEPR integrated project EPiQ ANR-22-PETQ-0007 part of Plan France 2030 and by ANR JCJC TCS-NISQ ANR-22-CE47-0004. Q-H.V was supported in part by the French ANR project CTQC (ANR-25-CE48-7926-01 CTQC).

In the previous version of this work, Theorem 5.3 was left as a geometric conjecture. We managed to prove this conjecture using tools from differential geometry and isoperimetric inequalities, with the assistance of `ChatGPT 5.1 Thinking`. The proof in this paper is fully written by the authors, yet the main idea is from

ChatGPT 5.1 Thinking.

2 Technical overview

(Unitarized) Common Haar function-like state oracles and PRFSGs. Our results are in a relativized world with (variants of) the common Haar function-like state (CHFS) oracles. The CHFS oracles with length ℓ are defined as follows: it is a family of unitaries $\{S_x\}_{x\in\{0,1\}^*}$ defined as follows:

$$S_x : \begin{cases} |0\rangle \to |\phi_x\rangle \\ |\phi_x\rangle \to |0\rangle \\ |\psi\rangle \to |\psi\rangle \quad \text{if } |\psi\rangle \notin \mathsf{span}(|0\rangle, |\phi_x\rangle), \end{cases}$$

where $|\phi_x\rangle$ is a predetermined Haar random state of length $\ell(|x|)$, with $|x|$ denoting the bit-length of x. This oracle is inspired by the reflection/swap oracles in [GZ25, BMM+25].

In this overview, we assume that the algorithm accesses the unitaries S_x one by one, and also assume that $\langle 0|\phi_x\rangle = 0$ for simplicity, so that S_x can be understood as a reflection

$$S_x = I - 2\,|\phi_x-\rangle\langle\phi_x-|\,,$$

where $|\phi-\rangle = \frac{|0\rangle-|\phi_x\rangle}{\sqrt{2}}$.

The construction of PRFSGs with the CHFS oracles is rather straightforward: the generation algorithm, on input (k, x) for key k and input x of length λ, outputs $|\phi_{k||x}\rangle$ by querying $S_{k||x}$, where $k||x$ is the concatenation of k and x. Note that the output length of the PRFSGs is $\ell(k||x)$. The security can be shown by the standard reduction to the unstructured search problem.

2.1 Separating QPRGs from short PRFSGs

We show that relative to *some* CHFS oracles with output length $\ell(n)$, together with the PSPACE oracle, PRFSGs with output length ℓ exist but QPRGs do not. But first, we discuss a technical hurdle we met.

Concentration inequality fails. The concentration inequality for Haar measure (see Theorem 3.1) is the most common tool currently used for oracle separations. However, when the oracle outputs only $\ell(n) = \Theta(\log n)$ qubits, the ambient dimension is too small (because $2^\ell = \mathsf{poly}(n)$) for these bounds to force the type of "almost-everywhere" behavior needed to rule out QPRGs.

We instead start from an *extreme* concentration phenomenon that any QPRG must satisfy. Consider a single-bit-output QPRG G^O relative to CHFS oracles O with negligible errors. For any fixed seed x, the output distribution of $G^O(x)$ must be almost deterministic for all but a negligible fraction of oracles O: equivalently, defining $f(O) = \Pr\left[G^O(x) \to 1\right]$, we must have $f(O) \approx 0$ or $f(O) \approx 1$ for almost all O. These are the two extreme points in the concentration inequality. A natural question is thus whether these two extreme points can be simultaneously concentrated.

The new Barrier theorem. We ask the following question: if $f(O) \in [0, 1]$ has substantial mass near both 0 and 1, what can we say about the geometry of the preimages $f^{-1}([0, \varepsilon])$ and $f^{-1}([1-\varepsilon, 1])$ in oracle space? If $G^O(x)$ for a fixed x can output both 0 and 1 with non-zero probability, both pre-image regions are large. We also expect the distance between the two pre-image regions to be large, since close oracles would likely induce close outputs. Our Barrier theorem (Theorem 5.3) asserts that under such conditions, the intermediate region $f^{-1}((\varepsilon, 1-\varepsilon))$ must itself be large:

Theorem 2.1 (Barrier theorem (Informal))**.** *Let X be the product space of pure quantum states with the corresponding product Haar measure σ. If S_0, S_1 are two measurable subsets of X such that $\sigma(S_0), \sigma(S_1) \geq A$, and if $d(S_0, S_1) \geq B$ for some distance d on X, then $\sigma(X \setminus (S_0 \cup S_1)) \geq cAB$ for a universal constant $c > 0$ not depending on X.*

Now we turn back to the QPRGs G^O with negligible pseudodeterminism error. Leveraging the Barrier theorem, our key observation is that with such strong pseudodeterminism, for a fixed seed x the output cannot meaningfully depend on the CHFS oracles, otherwise we would see a noticeable "intermediate" mass of oracles where the output is not deterministic. In more detail, consider the first output bit of G^O denoted by G_1^O and let $f(O) = \Pr_G[G_1^O(x) \to 1]$. The Barrier theorem rules out the case where $S_0 = f^{-1}([0, \varepsilon])$ and $S_1 = f^{-1}([1-\varepsilon, 1])$ are both large, because it leads that G^O is not pseudodeterministic on the too large barrier $X \setminus (S_0 \cup S_1)$. This leads to a dichotomous intuition: either

1. G_1^O is not pseudodeterministic, or
2. G_1^O is essentially constant independent of O.

That is, G_1^O must fail to achieve pseudodeterminism or security.

It turns out that this intuition is not quite true. One of the reasons is that the barrier theorem, even if the second statement miserably failed, shows that the first statement holds only for a mildly large fraction of O. A more correct dichotomy is as follows, considering all output bits. For a fixed security parameter, one of the following holds:

1. G^O is not pseudodeterministic on a (slightly) large fraction of O, or
2. G^O outputs a fixed value for a (very) large fraction of O.

We make an observation for each case. For the first case, we observe that the PRFSG construction $\mathsf{Gen} : (k, x) \to |\phi_{k||x}\rangle$ is still secure even if we choose O from a slightly smaller sub-distribution of the product Haar random distributions (at parameter λ). It means we may hope to choose O such that Gen is a secure PRFSG yet G is not pseudodeterministic.

For the second case, we observe that G^O can be estimated by $G^{O'}$ for randomly chosen O', in which case we know how to simulate without querying O using known techniques.[15] Let F be the simulated function, then given input y, determining if there is x such that $F(x)$ can be done in polynomial space (i.e., with the PSPACE oracle), breaking the pseudorandomness of G^O.

Based on this observation, we sample oracle O (at parameter λ) depending on which case occurs: If G_i^O is pseudodeterministic with high probability over O for all i, we just sample O from the product Haar random distribution, breaking the pseudorandomness. For the other case, the dichotomy says that a (slightly) large fraction of O makes G^O not pseudodeterministic, thus we sample O among one that makes G^O not pseudodeterministic yet Gen is secure, given that the fractions for G not being pseudodeterministic are larger than the fractions for insecure Gen. This is indeed possible by adjusting parameters.

This, however, only ensures that a single candidate G is insecure or not pseudodeterministic at a single parameter. Our final proof of Theorem 1.1 proceeds by diagonalization, considering all uniform oracle-aided QPRG candidates $\{G_j\}_{j\in\mathbb{N}}$ as follows.

1. Define $\{G'_j\}_{j\in\mathbb{N}}$ so that for each i there are infinitely many j's such that $G_i = G'_j$ holds. For example, one can consider $\{G_i\}_i$ as $G_1, G_1, G_2, G_1, G_2, G_3, \ldots$.
2. Choose a super-fast-growing sequence of security parameters $\{\lambda_i\}_{i\in\mathbb{N}}$ (e.g. $\lambda_{i+1} = 2^{\lambda_i^{200}}$) and blocks of oracle input lengths $I_i = [\log(\lambda_i), \lambda_i^{100}]$ that are disjoint.

   Intuitively, the properties of G'_i at parameter λ_i *essentially* depend only on the oracle with input lengths in I_i. This is because $G'_i(1^{\lambda_i}, \cdot)$ may never query too large input length $|x| > \lambda_i^{100}$, and for the query $|x| < \log(\lambda_i)$, it can be efficiently estimated by the adversary using tomography.
3. Fix an oracle $O_{<i}$ up to input length $< a_i$. Over the random choice of O with respect to the input length in I_i, the dichotomy holds: either $(G'_i)^O$ is not pseudodeterministic with high probability over O, or $(G'_i)^O$ can be simulated with high probability over O so that it is insecure. We sample O over the input length in I_i as above, according to which case occurs.

[15] This is the final goal of the concentration inequality used in [Kre21]!

4. Proceed to the next block after sampling O, by fixing $O < i + 1$.

The above procedure results in an oracle $(O, \mathbf{PSPACE})$ relative to which Gen is a secure PRFSG, while any QPRG candidate G must be infinitely often insecure or infinitely often non-pseudodeterministic. Note that in the actual proof, we need to consider each oracle QPT adversary $\mathcal{A}$ instead of the one-shot security argument of Gen, which further complicates the argument.

2.2 On separating PRUs from PRFSGs

We now give a brief overview on how we achieve Theorem 1.4. We consider the unitary CHFS oracles with output length $\ell(n) = n$. As discussed above, we can easily construct PRFSGs relative to this oracle, but breaking the PRU constructions is quite involved. We sketch the outline of the proof here.

Breaking PRUs without ancilla. To establish Theorem 1.4, we present an explicit attack for any PRU candidate without ancilla with respect to the CHFS oracle of length $\ell(n) = n$.

We consider the following simplified form of the PRU algorithm $\{G_k\}_{k\in\{0,1\}^*}$ on key $k \in \{0,1\}^\lambda$ and input state $|\psi\rangle$:

$$G_k : |\psi\rangle \mapsto U_T^{(k)} \cdot S_{x_T^{(k)}} \cdot U_{T-1}^{(k)} \cdot \ldots \cdot U_1^{(k)} \cdot S_{x_1^{(k)}} \cdot U_0^{(k)} |\psi\rangle ,$$

where $U_T^{(k)}, \ldots, U_0^{(k)}$ are some unitaries and $S_{x_T^{(k)}}, \ldots, S_{x_1^{(k)}}$ are the CHFS oracle queries. In the main body of the paper, we consider a more general form of G_k where queries may be in superposition or adaptive.

Our main observation is as follows: for a Haar random state $|\rho\rangle$ independently chosen from the oracle, the application of the reflection oracle does not change the state much, i.e.,

$$S_x |\rho\rangle \approx |\rho\rangle . \tag{1}$$

This is because the reflection S_x only makes a change on the tiny space spanned by $\{|\phi_x\rangle, |0\rangle\}$. Therefore, one may argue that

$$G_k |\rho\rangle \approx U_T^{(k)} \cdot U_{T-1}^{(k)} \cdot \ldots \cdot U_1^{(k)} \cdot U_0^{(k)} |\rho\rangle$$

because $|\rho_t\rangle := U_t^{(k)} \cdot \ldots \cdot U_0^{(k)} |\rho\rangle$ is a Haar random state independent of the oracle due to the invariant property of Haar measure. Unfortunately, this is not the case in general, as the loss in Eq. (1) is proportional to $1/2^{|x|}$, so we cannot ignore S_x for small $|x|$.

We instead learn all S_x to obtain S'_x for small $|x|$ using process tomography [HKOT23]. We define $\tilde{S}_x$ by S'_x for small $|x|$ and I for large $|x|$, and define

$$F_k : |\psi\rangle \mapsto U_T^{(k)} \cdot \tilde{S}_{x_T^{(k)}} \cdot U_{T-1}^{(k)} \cdot \ldots \cdot U_1^{(k)} \cdot \tilde{S}_{x_1^{(k)}} \cdot U_0^{(k)} |\psi\rangle$$

which now satisfies $F_k |\rho\rangle \approx G_k |\rho\rangle$.

Now we describe the adversary that given oracle V, distinguishes whether it is one of $\{G_k\}$ or a true Haar random unitary. The adversary first prepares $\Phi = (|\rho\rangle \otimes V |\rho\rangle)^{\otimes M}$ for some large M and Haar random state $|\rho\rangle$ (or a t-design for sufficiently large t) and defines:

P_k**:** on input $\Phi = (|\rho\rangle \otimes V |\rho\rangle)^{\otimes M}$, it applies $(F_k \otimes I)^{\otimes M}$, applies M swap tests on each copy; if sufficiently many copies pass the swap test, it returns 1. Otherwise, it returns 0.

We can show that P_k returns 1 if $V = G_k$ with high probability, but P_k almost always returns 0 if V is a Haar random unitary. This satisfies the setting where the quantum OR tester [HLM17] can be run with the **QPSPACE** oracle[16] as observed in [CCS25]. By augmenting our world with the **QPSPACE** oracle, we obtain a relativized world where PRFSGs exist but PRUs without ancilla and measurements do not, proving Theorem 1.4.

The attack even breaks the non-adaptive PRU security as V is only used to prepare Φ. Extending this to the quantum-accessible PRFSGs security requires considering the coherent version of CHFS oracles, which can be done similarly with some more computation.

[16]This oracle, roughly, takes a quantum state $|\phi\rangle$ and a succinct description of a quantum circuit C computable in polynomial space, and returns $C |\phi\rangle$. See Definition 3.9 for the formal definition.

2.3 On the length extension of PRSGs

We now give an overview on how to prove Theorem 1.5. We again consider the CHFS oracle with $\ell(|x|) = \lfloor \log |x| \rfloor$ together with the **QPSPACE** oracle. Here, we consider the classical-accessible isometry version: a family of isometries $\{O_x\}_{x\in\{0,1\}^*}$ where O_x takes input $|0\rangle$ and outputs an $\ell(|x|)$-qubit Haar random state $|\phi_x\rangle$. We do not allow querying the other input states.

Consider the following PRSG candidate that makes nonadaptive queries to the oracle and outputs on key k

$$\rho_k = U_k \left(\left| \phi_{x_1^{(k)}} \right\rangle \otimes \cdots \otimes \left| \phi_{x_t^{(k)}} \right\rangle \right), \tag{2}$$

where we assume that the parameter t and the lengths of $x_i^{(k)}$'s are all the same for different keys for simplicity in this overview. Here $\left|\phi_{x_1^{(k)}}\right\rangle,\ldots,\left|\phi_{x_t^{(k)}}\right\rangle$ are shorter PRSG outputs. We have that the state

$$U_k^\dagger \rho_k U_k = \left| \phi_{x_1^{(k)}} \right\rangle \otimes \cdots \otimes \left| \phi_{x_t^{(k)}} \right\rangle$$

is a product of many pure states. On the other hand, for a Haar random state $|\psi\rangle$, $\tilde{U}_k^\dagger |\psi\rangle$ definitely does not have such a product structure, as it is also Haar random by definition. Given the efficient product test algorithm [HM10], we can run the quantum OR tester with the **QPSPACE** oracle as in the separation between PRUs and PRFSGs.

The formal proof considers more general algorithms allowing partial traces with constant-size ancilla that can be reset to 0.

3 Preliminaries

Notations. We use $\lambda \in \mathbb{N}$ to denote the security parameter. For any $m \in \mathbb{N}$, we use the notation $[m]$ to refer to the set $\{1,\ldots,m\}$. For any finite set U, we write $x \leftarrow U$ to denote that x is sampled uniformly at random from U. For a distribution $\mathcal{D}$, $x \leftarrow \mathcal{D}$ denotes that x is sampled from $\mathcal{D}$. For a bit string $x \in \{0,1\}^*$, we denote its bit-length by $|x|$. We assume that all functions used to represent the lengths of the cryptographic primitives are QPT-computable. We assume the reader is familiar with the basics of quantum computation, and refer to [NC10] otherwise.

We will also use standard notations from quantum information and cryptography.

3.1 Quantum states, operations, and trace

A d-dimensional quantum state is a positive semi-definite Hermitian density matrix $\rho = \sum_{x\in[d]} p_x |\phi_x\rangle\langle\phi_x|$, where the pure states $|\phi_x\rangle\langle\phi_x|$ have trace one, and $p_1,\ldots,p_d$ is a probability distribution, i.e., $p_1,\ldots,p_d \geq 0$ and $p_1+\cdots+p_d = 1$. Pure states are the rank-1 quantum state that can be written as $|\phi\rangle\langle\phi|$. We sometimes write $|\phi\rangle$ or just ϕ to denote the pure state $|\phi\rangle\langle\phi|$ for simplicity. We can consider any positive semi-definite Hermitian matrix (with arbitrary unit trace) as an unnormalized quantum state, e.g., $\Pi\rho\Pi$ for some projection Π and quantum state ρ, and call them unnormalized states.

A quantum operation Φ is a completely positive and trace-nonincreasing operator, that can be represented by matrices $B_1,\ldots,B_k$ satisfying

$$I - \sum_{i=1}^{k} B_i^\dagger B_i \geq 0.$$

The matrices $B_1,\ldots,B_k$ are the Kraus operators of the operations, and with this notation, Φ maps a quantum state ρ to $\Phi(\rho) = \sum_{i=1}^k B_i \rho B_i^\dagger$. Quantum operations can represent unitary operations, projective measurements, or applying a projection Π. We write the composition of two quantum operations Φ, Ψ by $\Phi \circ \Psi$. For a unitary U, the corresponding operation is represented by $U(\rho) = U\rho U^\dagger$ or sometimes the calligraphic font $\mathcal{U}(\rho)$.

The trace norm of a Hermitian matrix A is defined by $\|A\|_1 := \sum_{i=1}^{d} |\lambda_i|$, where $\lambda_1, \ldots, \lambda_d$ are the eigenvalues of A. If A is positive semi-definite, we can write $\|A\|_1 = \mathrm{Tr}(A)$. This induces the *trace distance* $\|\rho - \sigma\|_{tr} = \frac{1}{2}\|\rho - \sigma\|_1$ between two (possibly unnormalized) mixed states, which forms a distance over (unnormalized) mixed states. A quantum operation Φ does not increase the trace norm. That is, for any Hermitian matrix A, it holds that $\|\Phi(A)\|_1 \leq \|A\|_1$. In particular, we have $\mathrm{Tr}(\Phi(A)) \leq \mathrm{Tr}(A)$ for any positive semi-definite matrix A. For any two (possibly unnormalized) states ρ, σ,

$$\|\Phi(\rho) - \Phi(\sigma)\|_{tr} = \frac{1}{2}\|\Phi(\rho - \sigma)\|_1 \leq \frac{1}{2}\|\rho - \sigma\|_1 = \|\rho - \sigma\|_{tr}. \tag{3}$$

For a positive semi-definite matrix A, it holds that

$$\mathrm{Tr}\left(A^2\right) \leq \mathrm{Tr}(A)^2. \tag{4}$$

We stress that most of the facts on the trace norm and distance also hold for unnormalized states, i.e., positive semi-definite Hermitian matrices.

3.2 Haar random states and unitaries

We write $\mathbb{S}(N)$ and $\mathbb{U}(N)$ to denote the set of N-dimensional pure quantum states and the group of $N \times N$ unitary matrices. We denote by σ_n (or sometimes $\mathcal{H}_{n(\lambda)}$) and μ_n the Haar distribution over n-qubit states and n-qubit unitaries, i.e., over $\mathbb{S}(2^n)$ and $\mathbb{U}(2^n)$, respectively. When the dimension is clear from the context, we drop the parameter and use σ or μ. The Frobenius norm $\|A\|_F$ of a matrix A is defined by $\sqrt{\mathrm{Tr}(A^\dagger A)}$.

Theorem 3.1 ([Mec19, Theorem 5.17])**.** *Let $n_1, \ldots, n_k \in \mathbb{N}$ and $\mu = \mu_{n_1} \times \cdots \times \mu_{n_k}$ be the product of Haar unitary measures over $X = \mathbb{U}(2^{n_1}) \times \cdots \times \mathbb{U}(2^{n_k})$. Suppose that $f : X \to \mathbb{R}$ is L-Lipschitz in the Frobenius norm. Let $N = \min(2^{n_1}, \ldots, 2^{n_k})$. For every $t > 0$, it holds that*

$$\Pr_{U \leftarrow \mu}\left[f(U) \geq \mathbb{E}_{V \leftarrow \mu}[f(V)] + t\right] \leq \exp\left(-\frac{(N-2)t^2}{24L^2}\right).$$

Corollary 3.2. *Let C^U be an m-query quantum oracle algorithm for the product of Haar random unitaries U chosen from X according to μ defined above. Let $g(U) := \Pr\left[1 \leftarrow C^U\right]$. Then it holds that*

$$\Pr_{U \leftarrow \mu}\left[g(U) \geq \mathbb{E}_{V \leftarrow \mu}[g(V)] + t\right] \leq \exp\left(-\frac{t^2(N-2)}{24m^2}\right).$$

Proof. In [Kre21], the following statement is shown.

Lemma 3.3 ([Kre21])**.** *Let A^U be a quantum algorithm that makes T queries to the unitary oracle U. Define $f(U) := \Pr\left[1 \leftarrow A^U\right]$. Then f is T-Lipschitz in the Frobenius norm, i.e., $|f(U) - f(V)| \leq T \cdot \|U - V\|_F$.*

This lemma ensures that C is m-Lipschitz, thus applying Theorem 3.1, we obtain the desired result. □

Lemma 3.4. *For any rank-D projection Π on m qubits for $m \geq n$,*

$$\mathbb{E}_{|\phi\rangle \leftarrow \sigma_n}\left\langle \phi, 0^{m-n}\right| \Pi \left|\phi, 0^{m-n}\right\rangle \leq \frac{D}{2^n}.$$

If $m = n$, the equality holds. In particular, for any n-qubit mixed state ρ, $\mathbb{E}_{|\phi\rangle \leftarrow \sigma_n}\langle\phi|\rho|\phi\rangle = \frac{1}{2^n}$.

Proof. We simply write 0 to denote 0^{m-n}. We can write $\mathbb{E}_{|\phi\rangle \leftarrow \sigma_n}\langle\phi, 0|\Pi|\phi, 0\rangle$ by

$$\mathbb{E}_{|\phi\rangle \leftarrow \sigma_n} \mathrm{Tr}(\Pi \cdot |\phi, 0\rangle\langle\phi, 0|) = \mathrm{Tr}\left(\Pi \cdot \frac{I \otimes |0\rangle\langle 0|}{2^n}\right) \leq \frac{1}{2^n}\mathrm{Tr}(\Pi) = \frac{D}{2^n},$$

where the last equality follows from the fact that $\mathrm{Tr}(\Pi) = \mathrm{rank}(\Pi)$. If $m = n$, the inequality is saturated. The last statement can be shown by writing $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ for $\sum_i p_i = 1$. □

3.3 Cryptographic primitives

We define cryptographic primitives relative to an oracle O.

Definition 3.5 (PRSGs)**.** *We say that an oracle QPT algorithm Gen^O is a secure pseudorandom state generator (PRSG) in the CHFS model if the following holds for some functions $\kappa, n : \mathbb{N} \to \mathbb{N}$ such that $\kappa = \omega(\log \lambda)$:*

- ***State Generation:*** *For any $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^{\kappa(\lambda)}$, the algorithm $\mathsf{Gen}^O(k)$ outputs an $n(\lambda)$-qubit state.*
- ***Pseudorandomness:*** *For any polynomial $t(\cdot)$ and any oracle QPT adversary $\mathcal{A}^O = \{\mathcal{A}^O_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a negligible function $\varepsilon(\cdot)$ such that for all $\lambda \in \mathbb{N}$:*

$$\left|\Pr_{k\leftarrow\{0,1\}^{\kappa(\lambda)}}\left[1 \leftarrow \mathcal{A}^O_\lambda(\mathsf{Gen}^O(k)^{\otimes t(\lambda)})\right] - \Pr_{|\psi\rangle\leftarrow\sigma_{n(\lambda)}}\left[1 \leftarrow \mathcal{A}^O_\lambda(|\psi\rangle^{\otimes t(\lambda)})\right]\right| \leq \varepsilon(\lambda).$$

We say that Gen^O is a $n(\lambda)$-PRSG to indicate that its output length is $n(\lambda)$. We further say that a PRSG is a short PRSG *when its output length is $\Theta(\log \lambda)$, and a* (long) PRSG *when its output length is $\omega(\log \lambda)$.*

From now on we will use PRSGs to refer to long PRSGs and short PRSGs for logarithmic output. We by default consider the adaptively-secure PRFSGs defined as follows.

Definition 3.6 (PRFSGs)**.** *We say that a QPT algorithm Gen^O is a secure pseudorandom function-like state generator (PRFSG) in the CHFS model if the following holds for some functions $\kappa, m, n : \mathbb{N} \to \mathbb{N}$ such that $\kappa, m = \omega(\log \lambda)$:*

- ***State Generation:*** *For any $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^{\kappa(\lambda)}$, the algorithm Gen^O_k takes as input $x \in \{0,1\}^{m(\lambda)}$ and outputs $n(\lambda)$-qubit (possibly mixed) state $\mathsf{Gen}^O_k(x)$ stored in a new register.*
- ***Pseudorandomness:*** *For any oracle QPT adversary $\mathcal{A}^O = \{\mathcal{A}^O_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a negligible function $\varepsilon(\cdot)$ such that for all $\lambda \in \mathbb{N}$:*

$$\left|\Pr_{k\leftarrow\{0,1\}^{\kappa(\lambda)}}\left[1 \leftarrow \mathcal{A}^{O,\mathsf{Gen}^O(k,\cdot)}_\lambda\right] - \Pr_{G_{\mathsf{Haar}}}\left[1 \leftarrow \mathcal{A}^{O,G_{\mathsf{Haar}}(\cdot)}_\lambda\right]\right| \leq \varepsilon(\lambda),$$

where $G_{\mathsf{Haar}}(\cdot)$ on input $x \in \{0,1\}^{m(\lambda)}$, output $|\psi_x\rangle$ stored in a new register, where, for every $x \in \{0,1\}^{m(\lambda)}$, $|\psi_x\rangle \leftarrow \mathcal{H}_{n(\lambda)}$.

When the adversary always measures the input register before making queries to Gen^O_k or G_{Haar}, we say that Gen^O_k is classical-accessible. Otherwise, we say that it is quantum-accessible.

We say that Gen is a $(\kappa(\lambda), m(\lambda), n(\lambda))$-PRFSG to indicate that its key length is $\kappa(\lambda)$, its input length is $m(\lambda)$, and its output length is $n(\lambda)$. We say that a PRFSG is a short PRFSG *when $n = \Theta(\log \lambda)$, and a* (long) PRFSG *when $n = \omega(\log \lambda)$.*

For pseudorandom unitaries, we only consider the super-logarithmic output length and without inverse oracle access. Unlike [JLS18], we allow PRUs to not be unitary.

Definition 3.7 (PRUs)**.** *We say that an oracle QPT algorithm G^O is a pseudorandom unitary in the CHFS model if the following holds for some $n : \mathbb{N} \to \mathbb{N}$ such that $n = \omega(\log \lambda)$:*

- ***Quantum operation:*** *For any $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^\lambda$, G^O_k takes as input an $n(\lambda)$-qubit (mixed) state ρ and outputs an $n(\lambda)$-qubit state $G^O_k(\rho)$.*
- ***Pseudorandomness****: For any oracle QPT adversary $\mathcal{A}^O = \{\mathcal{A}^O_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a negligible function ε such that for all $\lambda \in \mathbb{N}$,*

$$\left|\Pr_{k\leftarrow\{0,1\}^\lambda}\left[1 \leftarrow \mathcal{A}^{O,G^O_k}_\lambda\right] - \Pr_{\mathcal{U}\leftarrow\mu_{n(\lambda)}}\left[1 \leftarrow \mathcal{A}^{O,\mathcal{U}}_\lambda\right]\right| \leq \varepsilon(\lambda).$$

When the adversary makes non-adaptive queries to G^O and $\mathcal{U}$, we say that G is non-adaptively secure.

We also define quantum pseudorandom generators (QPRGs), which are algorithms whose output is indistinguishable from random, and is always the same with probability negligibly close to one.

Definition 3.8 (QPRGs)**.** *We say that an oracle QPT algorithm F^O that outputs an $m(\lambda)$-bit classical string on $n(\lambda)$-bit input is a* quantum pseudorandom generator *(QPRG) with $(1-\varepsilon)$-pseudodeterminism if the following conditions hold for a function ε. If ε is negligible, we just call F a QPRG.*

- $(1-\varepsilon)$***-Pseudodeterminism****. For every $\lambda \in \mathbb{N}$, the following holds:*
  1. *There exists a set $K_\lambda \subseteq \{0,1\}^{n(\lambda)}$ of "good seeds" such that*
  $$\Pr\left[x \in K_\lambda : x \leftarrow \{0,1\}^{n(\lambda)}\right] \geq 1-\varepsilon.$$
  2. *There exists a deterministic function $f_\lambda : \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}$ such that for every $x \in K_\lambda$, it holds that*
  $$\Pr\left[F^O(x) = f_\lambda(x)\right] \geq 1-\varepsilon,$$
  where the probability is over the randomness of F.
- ***Security****. For any oracle QPT algorithm $\mathcal{A}^O = \{\mathcal{A}^O_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a negligible function ε such that*
$$\left|\Pr_{y\leftarrow\{0,1\}^{m(\lambda)}}\left[1 \leftarrow \mathcal{A}^O_\lambda(y)\right] - \Pr_{x\leftarrow\{0,1\}^{n(\lambda)}}\left[1 \leftarrow \mathcal{A}^O_\lambda(F^O(x))\right]\right| \leq \varepsilon(\lambda),$$
where the probability is over the randomness of F and $\mathcal{A}_\lambda$.
- ***Length extension.*** *$n(\lambda) < m(\lambda)$ holds for all $\lambda \in \mathbb{N}$.*

3.4 QPSPACE oracle

We recall the definition of the QPSPACE oracle that implements the arbitrary unitary operation described by polynomial size input [CCS25, BMM+25].

Definition 3.9 (**QPSPACE** Oracle)**.** *The* unitary *QPSPACE machine oracle, denoted by* **QPSPACE***, is defined as follows: it takes a pair (ρ, M, t, b) of an ℓ-qubit quantum state ρ, a classical Turing machine M, and an integer $t \in \mathbb{N}$, and $b \in \{0,1\}$ denotes the abort bit. The oracle runs M for t steps to obtain the description of a unitary quantum circuit C that operates on ℓ qubits; if M does not terminate after t steps or the output is not described as above, the oracle flips the abort register and halts. Otherwise, the oracle applies C on ρ and returns the output quantum state without measurement.*

The quantum access to the QPSPACE oracle is done by allowing coherent (M, t). For any unitary quantum circuit C that is output by a machine M after t steps, there is a QPT algorithm with **QPSPACE** oracle that implements $C^{-1}(\rho)$ on input ρ [BMM+25, Proposition 3.5].

3.5 State property tests

Swap test

We review the basic results of the swap test, which can be used to test the purity of a state. We provide some lemmas about the swap test on a state that is close to pure states, which are essential to obtain our results.

For two quantum states σ, ρ stored in two different registers $\mathbf{A}, \mathbf{B}$, the swap test is executed on the registers $\mathbf{A}, \mathbf{B}$ and a control register $\mathbf{C}$ initialized to $|1\rangle\langle 1|$. It applies Hadamard on $\mathbf{C}$, swaps $\mathbf{A}$ and $\mathbf{B}$ conditioned on $\mathbf{C}$, and measures $\mathbf{C}$ on the Hadamard basis.

Lemma 3.10 (Swap test)**.** *The swap test on input (σ, ρ) outputs 1 with probability*

$$\frac{1 + \mathrm{Tr}(\rho\sigma)}{2},$$

in which case we say that it passes the swap test. For pure states $|\sigma\rangle, |\rho\rangle$, it equals $\frac{1+|\langle\rho|\sigma\rangle|^2}{2}$.

When $\sigma = \rho$, we sometimes call it a purity test on ρ, which outputs 1 with certainty if and only if ρ is a pure state.

Lemma 3.11. *Suppose that $\mathrm{Tr}(\rho^2) \leq 1 - 1/T$ for some state ρ and $T \in \mathbb{N}$. Let $\lambda \in \mathbb{N}$. If we run the purity test $16T\lambda$ times on ρ, then the probability that at least 4λ tests fail among $16T\lambda$ is at least $1 - 2^{-\lambda}$.*

Proof. Note that each test succeeds with probability $(1 + \mathrm{Tr}(\rho^2))/2 \leq 1 - 1/2T$, and is independent of each other. Applying Chernoff's inequality (Lemma 3.20) for $\delta = 1/2$, we obtain the desired result. □

Product test

We first recall the product test to determine whether an n-partite state $|\phi\rangle$ is a product state or far from any product state from [HM10], then give a bound on the success of the product test on Haar-random states.

Lemma 3.12 ([HM10, Lemma 3], Product test for mixed states)**.** *Let $m \in \mathbb{N}$ and $d_1, \ldots, d_m$ be the local dimensions of a n-qubit system, i.e. $\prod_{i\in[m]} d_i = 2^n$. Let ρ be a mixed state of n-qubits and for every $S \subseteq [m]$, denote by ρ_S the state after tracing out the subsystem $\overline{S} := [m] \setminus S$. Let $\mathcal{A}_{\mathrm{PTEST}}$ denote the algorithm that, given two copies of ρ, performs the swap test on each of the m pairs of corresponding subsystems of the two copies of ρ, and that outputs 1 if all the tests succeed, and 0 otherwise. Then, the probability that the algorithm $\mathcal{A}_{\mathrm{PTEST}}$ outputs 1 when applied to two copies of ρ is equal to*

$$\Pr\left(1 \leftarrow \mathcal{A}_{\mathrm{PTEST}}(|\phi\rangle^{\otimes 2})\right) = \frac{1}{2^m} \sum_{S\subseteq[m]} \mathrm{Tr}\left[\rho_S^2\right].$$

For Haar-random states, the above formula is explicitly calculated for any partition $S \cup \overline{S}$ of $[m]$ by [Lub78]:

$$\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma} \mathrm{Tr}\left[\rho_S^2\right] = \frac{d_S + d_{\overline{S}}}{d_S \cdot d_{\overline{S}} + 1}.$$

As a consequence, we have the following bound for the success of the product test on Haar-random states.

Lemma 3.13 (Product test for Haar-random states)**.** *Let $m \in \mathbb{N}$ and $\{d_i\}_{i\in[m]}$ be the local dimensions of a n-qubit system, i.e. $\prod_{i\in[m]} d_i = 2^n$. Then, the probability that the algorithm $\mathcal{A}_{\mathrm{PTEST}}$ outputs 1 when applied to two copies of a n-qubit Haar-random state $|\psi\rangle$ satisfies:*

$$\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma} \Pr\left(1 \leftarrow \mathcal{A}_{\mathrm{PTEST}}(|\psi\rangle^{\otimes 2})\right) \leq 2\left(\frac{3}{4}\right)^m.$$

Proof. For every partition $S \cup \overline{S}$ of $[m]$, the local dimension of each partition is given by $d_S = \prod_{i\in S} d_i$.

$$\begin{aligned}
\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma} \Pr\left(1 \leftarrow \mathcal{A}_{\mathrm{PTEST}}(|\psi\rangle^{\otimes 2})\right) &= \mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\sigma} \left[\frac{1}{2^m} \sum_{S\subseteq[m]} \mathrm{Tr}\left[\rho_S^2\right]\right] \\
&= \frac{1}{2^m} \sum_{S\subseteq[m]} \frac{d_S + d_{\overline{S}}}{d_S \cdot d_{\overline{S}} + 1} \leq \frac{1}{2^m} \sum_{S\subseteq[m]} \frac{d_S + d_{\overline{S}}}{d_S \cdot d_{\overline{S}}} \\
&= \frac{1}{2^m} \left(\sum_{S\subseteq[m]} \frac{1}{d_S} + \frac{1}{d_{\overline{S}}}\right) = \frac{2}{2^m} \left(\sum_{S\subseteq[m]} \frac{1}{d_S}\right)
\end{aligned}$$

$$= \frac{2}{2^m} \prod_{i\in[m]} \left(1 + \frac{1}{d_i}\right) \leq \frac{2}{2^m} \prod_{i=1}^{m} \left(\frac{3}{2}\right) = 2\left(\frac{3}{4}\right)^m,$$

where we use the fact that each $d_i \geq 2$ to obtain the last inequality. □

We also need the following simple extension of the Haar product-test bound, in which the Haar-random input is restricted to a subspace obtained by appending zero ancillas.

Lemma 3.14. *Let $\mathcal{H}_0$ be a Hilbert space of dimension $d = 2^{n-t}$, and let $\mathcal{H} = \mathcal{H}_0 \otimes \mathbb{C}^{2^t}$, so that $\dim(\mathcal{H}) = D = 2^n$ for some $0 \leq t < n$. Let P be any quantum protocol that takes two copies of a pure state in $\mathcal{H}$ and outputs a bit. Suppose that*

$$\Pr_{|\phi\rangle\leftarrow\sigma_{\mathcal{H}}} \left[P\left(|\phi\rangle^{\otimes 2}\right) = 1\right] \geq 1 - \varepsilon.$$

Then

$$\Pr_{|\psi\rangle\leftarrow\sigma_{\mathcal{H}_0}} \left[P\left(\left(|\psi\rangle \otimes |0^t\rangle\right)^{\otimes 2}\right) = 1\right] \geq 1 - 4^t\varepsilon.$$

Proof. Since P is a two-copy quantum protocol with binary output, there exists a POVM element $0 \leq M \leq I$ on $\mathcal{H}^{\otimes 2}$ corresponding to the event that P outputs 0 such that, for every pure state $|\eta\rangle \in \mathcal{H}$,

$$\Pr\left[P\left(|\eta\rangle^{\otimes 2}\right) = 0\right] = \mathrm{Tr}\left(M\,|\eta\rangle\langle\eta|^{\otimes 2}\right).$$

Let $\Omega_{\mathcal{H}} \coloneqq \mathbb{E}_{|\phi\rangle\leftarrow\sigma_{\mathcal{H}}}\left[|\phi\rangle\langle\phi|^{\otimes 2}\right]$. By the Haar second-moment formula [Mel24],

$$\Omega_{\mathcal{H}} = \frac{\Pi^{\mathcal{H}}_{\mathrm{sym}}}{\binom{D+1}{2}} = \frac{2}{D(D+1)}\Pi^{\mathcal{H}}_{\mathrm{sym}},$$

where $\Pi^{\mathcal{H}}_{\mathrm{sym}}$ is the projector onto the symmetric subspace of $\mathcal{H}^{\otimes 2}$.

Similarly, define the embedded subspace $W \coloneqq \mathcal{H}_0 \otimes |0^t\rangle \subseteq \mathcal{H}$ and let $\Omega_W \coloneqq \mathbb{E}_{|\psi\rangle\leftarrow\sigma_W}\left[|\psi\rangle\langle\psi|^{\otimes 2}\right]$. Again by the Haar second-moment formula, now applied inside the subspace W,

$$\Omega_W = \frac{\Pi^{W}_{\mathrm{sym}}}{\binom{d+1}{2}} = \frac{2}{d(d+1)}\Pi^{W}_{\mathrm{sym}},$$

where Π^{W}_{sym} is the projector onto the symmetric subspace of $W^{\otimes 2}$.

Since $W \subseteq \mathcal{H}$, the symmetric subspace of $W^{\otimes 2}$ is a subspace of the symmetric subspace of $\mathcal{H}^{\otimes 2}$. Hence $\Pi^{W}_{\mathrm{sym}} \leq \Pi^{\mathcal{H}}_{\mathrm{sym}}$. Therefore,

$$\Omega_W = \frac{2}{d(d+1)}\Pi^{W}_{\mathrm{sym}} \leq \frac{2}{d(d+1)}\Pi^{\mathcal{H}}_{\mathrm{sym}} = \frac{D(D+1)}{d(d+1)}\Omega_{\mathcal{H}}.$$

By assumption,

$$\mathrm{Tr}(M\Omega_{\mathcal{H}}) = \Pr_{|\phi\rangle\leftarrow\sigma_{\mathcal{H}}}\left[P\left(|\phi\rangle^{\otimes 2}\right) = 0\right] \leq \varepsilon.$$

Thus,

$$\begin{aligned}
\Pr_{|\psi\rangle\leftarrow\sigma_{\mathcal{H}_0}}\left[P\left(\left(|\psi\rangle \otimes |0^t\rangle\right)^{\otimes 2}\right) = 0\right] &= \mathrm{Tr}(M\Omega_W) \\
&\leq \frac{D(D+1)}{d(d+1)}\,\mathrm{Tr}(M\Omega_{\mathcal{H}}) \\
&\leq \frac{D(D+1)}{d(d+1)}\varepsilon \\
&\leq \left(\frac{D}{d}\right)^2 \varepsilon = 4^t\varepsilon.
\end{aligned}$$

□

3.6 Quantum OR lemma

Lemma 3.15 ([HLM17, Corollary 3.1], Quantum OR lemma)**.** *Let $\{\Pi_i\}_{i\in[N]}$ be binary-valued POVMs. Let $0<\varepsilon<1/2$ and $\delta>0$. Let Ψ be a quantum state such that either*

i) there exists $i\in[N]$ such that $\mathrm{Tr}[\Pi_i\Psi]\geq 1-\varepsilon$, or

ii) for all $i\in[N]$, $\mathrm{Tr}[\Pi_i\Psi]\leq\delta$.

Then, there is a quantum circuit C, called "OR tester", such that measuring the first qubit in case i) yields

$$\Pr(1\leftarrow C(\Psi))\geq\frac{(1-\varepsilon)^2}{7},$$

and in case ii),

$$\Pr(1\leftarrow C(\Psi))\leq 4N\delta.$$

Moreover, the circuit C can be implemented by a unitary quantum poly-space machine as long as each POVM Π_i can be implemented by a quantum poly-space machine and the set of measurements has a concise polynomial description. In other words, the quantum OR tester can be executed by a **QPSPACE***-aided BQP algorithm, where the oracle* **QPSPACE** *is defined in Definition 3.9.*

Remark 3.16. "Moreover" part of the above theorem for the projective measurements is shown in [CCS25, Appendix A], and the extension to the POVMs is observed in [BMM+25, Lemma 5.2].

3.7 Useful lemmas

Lemma 3.17 (Almost as good as new lemma [Aar04, Aar16])**.** *Let $\mathcal{M}=(\Pi_0,\Pi_1)$ be a binary measurement that acts as $\mathcal{M}(\rho)=\Pi_0\rho\Pi_0+\Pi_1\rho\Pi_1$. If $\mathrm{Tr}[\Pi_0\rho]\geq 1-\varepsilon$ for $\varepsilon>0$, then it holds that $\|\rho-\mathcal{M}(\rho)\|_{tr}\leq\sqrt{\varepsilon}$.*

Corollary 3.18. *In the same setting, $\|\rho-\Pi_0\rho\Pi_0\|_{tr}\leq\varepsilon+\sqrt{\varepsilon}\leq 2\sqrt{\varepsilon}$.*

Proof. We have $\|\mathcal{M}(\rho)-\Pi_0\rho\Pi_0\|_{tr}=\|\Pi_1\rho\Pi_1\|_{tr}\leq\varepsilon$, which gives the result. □

Norms and Process tomography.

For a matrix M, the operator norm is defined by

$$\|M\|_{op}:=\sup_{\||\phi\rangle\|_2=1}\|M|\phi\rangle\|_2,$$

which satisfies $\|M+N\|_{op}=\max(\|M\|_{op},\|N\|_{op})$ if M and N act on the orthogonal space. In particular, for $M=\sum_x|x\rangle\langle x|\otimes M_x$, it holds that

$$\|M\|_{op}=\max_x\|M_x\|_{op}. \tag{5}$$

The diamond norm of an operator A, denoted by $\|A\|_\diamond$, is defined by:

$$\|A(\cdot)\|_\diamond:=\sup_{\mathrm{Tr}(\rho)=1,\rho\geq 0}\|A\otimes I(\rho)\|_1,$$

where I denotes the identity acting with the same dimension as A. We sometimes omit $(\cdot)$ if it is clear from the context. We use the following fact about the diamond norm: for quantum channels A,B and a density matrix ρ, it holds that

$$\|A\otimes I(\rho)-B\otimes I(\rho)\|_{tr}\leq\frac{1}{2}\|A(\cdot)-B(\cdot)\|_\diamond.$$

We will also use the fact that for unitaries U,V and the corresponding channels $\mathcal{U},\mathcal{V}$, it holds that

$$\|\mathcal{U}(\cdot)-\mathcal{V}(\cdot)\|_\diamond\leq 2\|U-V\|_{op} \tag{6}$$

for the operator norm $\|\cdot\|_{op}$ because

$$\begin{aligned}\|\mathcal{U}(\cdot)-\mathcal{V}(\cdot)\|_\diamond &= \sup_{\mathrm{Tr}(\rho)=1,\rho\geq 0}\|(U\otimes I)(\rho)(U^\dagger\otimes I)-(V\otimes I)(\rho)(V^\dagger\otimes I)\|_1\\ &\leq \sup_{\mathrm{Tr}(\rho)=1,\rho\geq 0}\|(U-V)\otimes I\|_{op}\|\rho\|_1\|U^\dagger\otimes I\|_{op}+\|V\otimes I\|_{op}\|\rho\|_1\|(U^\dagger-V^\dagger)\otimes I\|_{op}\\ &\leq 2\|U-V\|_{op}\end{aligned}$$

where we use $\|A\rho B\|_1\leq\|A\|_{op}\|\rho\|_1\|B\|_{op}$.

Theorem 3.19 ([HKOT23])**.** *There exists a quantum algorithm* Tom *that, given black-box access to a unitary Z acting on the d-dimensional space, satisfies the following for any input $\varepsilon,\delta\in(0,1)$:*

Accuracy: *It outputs a classical description of a unitary Z such that*

$$\Pr_{Z'\leftarrow\mathtt{Tom}}[\|\mathcal{Z}(\cdot)-\mathcal{Z}'(\cdot)\|_\diamond\leq\varepsilon]\geq 1-\delta.$$

Efficiency: *It makes $O\left(\frac{d^2}{\varepsilon}\log\frac{1}{\delta}\right)$ queries to Z, and takes $\mathsf{poly}(d,\frac{1}{\varepsilon},\log\frac{1}{\delta})$ time.*

Chernoff bounds.

We use the following concentration inequalities.

Lemma 3.20 (Multiplicative Chernoff bound)**.** *Let $X_1,\ldots,X_n$ be some independent random variables over $\{0,1\}$. Let $X=\sum_{i=1}^n X_i$ and $\mu=\mathbb{E}[X]$. It holds that*

- $\Pr[X\geq(1+\delta)\mu]\leq\exp\left(-\frac{\mu\delta^2}{2+\delta}\right)$ *for $\delta\geq 0$, and*
- $\Pr[X\leq(1-\delta)\mu]\leq\exp\left(-\frac{\mu\delta^2}{2}\right)$ *for $0<\delta<1$.*

4 Common Haar Function-like State Oracles

4.1 CHFS oracles and unitarization

We first recall the definition of *swap (or reflection) oracles*.

Definition 4.1. *For a n-qubit pure quantum state $|\phi\rangle$, the swap (or reflection) unitary is defined by*

$$S_{|\phi\rangle}:=|0^n\rangle\langle\phi|+|\phi\rangle\langle 0^n|+I_\perp=I-2\left|\phi-\right\rangle\left\langle\phi-\right|,$$

where we assume w.l.o.g. that $|\phi\rangle$ is orthogonal to $|0^n\rangle$, since if not, we can always append a single $|1\rangle$ to it in order to make it orthogonal. Here, $I_\perp$ is the identity on the subspace orthogonal to $\mathsf{span}\{|0^n\rangle,|\phi\rangle\}$ and $|\phi-\rangle=\frac{|0^n\rangle-|\phi\rangle}{\sqrt{2}}$.

The last equality implies that $S_{|\phi\rangle}$ is actually the reflection unitary with respect to $|\phi-\rangle$.

We proceed to define the length-ℓ common Haar-random function-like state (CHFS) oracle and its "unitarized" oracle. We fix a (QPT-computable) function $\ell:\mathbb{N}\to\mathbb{N}$ representing the output length for each oracle, where we typically consider $\ell(\lambda)=\Theta(\log\lambda)$ or $\ell(\lambda)=\lambda$. We define two versions of the CHFS oracles as follows.

Definition 4.2 (The isometry CHFS oracle)**.** *We denote by $\mathcal{O}_\ell$ the distribution over the family of isometry oracles where*

- **Randomness:** *Choose a $\ell(|x|)$-qubit Haar random quantum state $|\phi_x\rangle$ for each $x \in \{0,1\}^*$ and define $\Phi = \{|\phi_x\rangle\}_{x\in\{0,1\}^*}$.*
- **Setup:** *A family of oracles $O^\Phi = (O_x^\Phi)_{x\in\{0,1\}^*} \leftarrow \mathcal{O}_\ell$ is chosen by randomly sampling Φ, where $O_x^\Phi := |\phi_x\rangle\langle 0|$ denotes the isometry operator. Here $|0\rangle$ denotes the trivial quantum state of dimension 1.*
- **Query:** *It takes a quantum state $\rho_{\mathbf{XZ}}$ as input and applies the isometry*

$$O^\Phi := \sum_{x\in\{0,1\}^{|\mathbf{X}|}} |x\rangle\langle x|_{\mathbf{X}} \otimes O_x^\Phi = \sum_{x\in\{0,1\}^{|\mathbf{X}|}} |x\rangle\langle x|_{\mathbf{X}} \otimes |\phi_x\rangle_{\mathbf{Y}} \langle 0|,$$

on $\rho_{\mathbf{XZ}}$, where $\mathbf{Y}$ denotes a new $\ell(|\mathbf{X}|)$-qubit register, i.e., appending a new register $\mathbf{Y}$.

We say the CHFS oracle is classical-accessible *if the register $\mathbf{X}$ must always be measured in the computational basis before applying the query. Otherwise, we call the oracle* quantum-accessible.

Definition 4.3 (The unitarized CHFS oracle)**.** *We denote by $\mathcal{S}_\ell$ the distribution over the family of unitary oracles where*

- **Randomness:** *Choose a $\ell(|x|)$-qubit Haar random quantum state $|\phi_x\rangle$ for each $x \in \{0,1\}^*$ and define $\Phi = \{|\phi_x\rangle|1\rangle\}_{x\in\{0,1\}^*}$.*[17]
- **Setup:** *A family of oracles $S^\Phi = (S_x^\Phi)_{x\in\{0,1\}^*} \leftarrow \mathcal{S}_\ell$ is chosen by randomly sampling Φ, where $S_x^\Phi := S_{|\phi_x\rangle}$ denotes the reflection operator as defined in Definition 4.1.*
- **Query:** *It takes a quantum state $\rho_{\mathbf{XYZ}}$ as input such that $|\mathbf{Y}| = \ell(|\mathbf{X}|) + 1$ and applies the unitary*

$$S^\Phi := \sum_{x\in\{0,1\}^{|\mathbf{X}|}} |x\rangle\langle x|_{\mathbf{X}} \otimes S_x^\Phi = \sum_{x\in\{0,1\}^{|\mathbf{X}|}} |x\rangle\langle x|_{\mathbf{X}} \otimes S_{|\phi_x\rangle},$$

on $\rho_{\mathbf{XYZ}}$, where S_x is applied on the register $\mathbf{Y}$.

The classical-accessible and quantum-accessible unitarized CHFS oracles are defined analogously.

The *(length-ℓ) CHFS model* is defined as follows. The randomness Φ is chosen as an initialization. We note that the sets of randomness Φ used to define the isometry and unitarized CHFS oracles are the same. We omit the superscript Φ if the context makes it clear. Then, all parties have oracle access to the CHFS oracle $O = O^\Phi$ or $S = S^\Phi$. We call this the log-length CHFS model for $\ell(\lambda) = \mathcal{O}(\log\lambda)$, and the standard CHFS model for $\ell(\lambda) = \omega\log\lambda$.[18]

We call it the *state* (or *isometry*) CHFS model when the oracle is O^Φ, and the *unitary* (or *swap/reflection*) CHFS model when the oracle is S^Φ. We, however, occasionally use O to denote the (any) CHFS oracle if the context is clear.

4.2 Construction of PRFSGs in the CHFS model

We show that PRFSGs with output length ℓ exist in the length-ℓ CHFS model. Again, we stress that the PRFSGs are adaptively-secure by default.

Theorem 4.4. *Quantum-accessible (resp. classical-accessible) (κ, m, L)-PRFSGs exist in the length-ℓ quantum-accessible (resp. classical-accessible) CHFS model with probability 1 for any key size $\kappa = \omega(\log\lambda)$ and input size $m = \mathsf{poly}(\lambda)$ and $L = \ell(\kappa + m)$, regardless of the choice of unitary or isometry models. The same statement even holds relative to the* **PSPACE** *and* **QPSPACE** *oracles.*

[17] Here we explicitly append $|1\rangle$ to make the unitary CHFS oracle well-defined w.r.t. Definition 4.1. We occasionally omit $|1\rangle$ if there is no confusion.

[18] We usually consider the standard CHFS model with $\ell(\lambda) = \lambda$ for simplicity.

Proof. We define the following (κ, m, L)-PRFSGs. We explain the construction in the isometry CHFS model, but modifying it to the unitary CHFS model is obvious.

$\mathsf{Gen}^O(k,\cdot)$**:** On the m-qubit input register $\mathbf{X}$, it applies the map

$$|x\rangle_{\mathbf{X}} \to |x\rangle_{\mathbf{X}} \otimes |\phi_{k,x}\rangle\,.$$

This is done by, on input $\rho_{\mathbf{XZ}}$, appending the κ-qubit register $|k\rangle_{\mathbf{K}}$ and making a query to the oracle O^Φ on the register $\mathbf{KX}$ and discarding the registers $\mathbf{K}$.

We have that $|k|+|x| = m+\kappa = \mathsf{poly}(\lambda)$ thus Gen can be implemented by a BQP algorithm with a single query to the CHFS oracle with $m+\kappa$ length input.

We claim that this construction is a secure PRFSG. More precisely, we prove the following statement: For any algorithm A that makes q queries, define $\mathsf{Adv}(A,O)$ by

$$\Pr\left[A^{\mathsf{Gen}(k,\cdot),O}(1^\lambda)\to 1\right] - \Pr\left[A^{G_{\mathsf{Haar}}(\cdot),O}(1^\lambda)\to 1\right], \tag{7}$$

where $G_{\mathsf{Haar}}(\cdot)$ is sampled independently of O, and $G_{\mathsf{Haar}}(x)$ outputs an $L=\ell(\kappa+m)$-qubit Haar random state $|\psi_x\rangle$. Then, it holds that

$$\mathbb{E}_O\left[\mathsf{Adv}(A,O)^2\right] = \mathcal{O}\left(\frac{q^2}{2^\kappa}\right).$$

We prove this using the Yao's indistinguishability to unpredictability [Yao82] following [KQST23, Corollary 32]. Formally, we consider a quantum oracle algorithm B^s for $s\in\{0,1\}^{2^\kappa}$, whose goal is to determine whether $s=0^{2^\kappa}$ or $s=e_k$ for a uniformly random $k\in\{0,1\}^\kappa$.

The algorithm B samples its own CHFS oracle O', together with two independent instances $G_{\mathsf{Haar}}(\cdot)$ and $G'_{\mathsf{Haar}}(\cdot)$ as defined in Definition 3.6. For $z=(k',x)\in\{0,1\}^{\kappa+m}$, queries to the simulated CHFS oracle are answered using $G_{\mathsf{Haar}}(x)$ if $s_{k'}=1$, and using the state originally sampled for $O'(z)$ if $s_{k'}=0$. Queries on all other input lengths are answered using O' without modification. We denote the resulting CHFS oracle by O'_s. Each query to O'_s can be processed coherently using a constant number of queries to s.

We identify an output bit $b\in\{0,1\}$ with $2b-1\in\{-1,1\}$. The algorithm B proceeds as follows using A twice.

1. It runs $A^{G_{\mathsf{Haar}}(\cdot),O'_s}(1^\lambda)$, and denotes its output by $a\in\{-1,1\}$.
2. It samples $c\leftarrow\{-1,1\}$ uniformly.
3. If $c=1$, it samples $k'\leftarrow\{0,1\}^\kappa$, runs $A^{\mathsf{Gen}(k',\cdot),O'_s}(1^\lambda)$, where Gen is evaluated relative to O'_s, and denotes its output by $d\in\{-1,1\}$. If $c=-1$, it runs $A^{G'_{\mathsf{Haar}}(\cdot),O'_s}(1^\lambda)$, and again denotes its output by $d\in\{-1,1\}$.
4. It outputs 1 if $a=cd$, and outputs 0 otherwise.

The two executions of A use independent randomness.

For a fixed CHFS oracle O, let $p(O)$ denote the second probability in Eq. (7), and let $\Delta(O) := \mathsf{Adv}(A,O)$. Conditioned on the simulated CHFS oracle being O, the second execution of A satisfies

$$\mathbb{E}[cd\mid O] = \frac{1}{2}\left(2\big(p(O)+\Delta(O)\big)-1\right) - \frac{1}{2}\left(2p(O)-1\right) = \Delta(O).$$

If $s=0^{2^\kappa}$, then $G_{\mathsf{Haar}}(\cdot)$ is independent of O'_s, and the first execution of A corresponds to the ideal experiment. Moreover, O'_s has the same distribution as a CHFS oracle. It follows that

$$\Pr\left[B^{0^{2^\kappa}}\to 1\right] = \frac{1}{2} + \frac{1}{2}\mathbb{E}_O\left[\left(2p(O)-1\right)\Delta(O)\right].$$

On the other hand, if $s = e_k$, then $G_{\mathsf{Haar}}(x)$ is the state corresponding to (k, x) in O'_s for every x. Thus, for a uniformly random k, the first execution of A corresponds to the real experiment. The oracle O'_s still has the same distribution as a CHFS oracle, and k is independent of O'_s. Therefore,

$$\Pr_k[B^{e_k} \to 1] = \frac{1}{2} + \frac{1}{2}\mathbb{E}_O\left[\left(2\left(p(O) + \Delta(O)\right) - 1\right)\Delta(O)\right].$$

Consequently,

$$\Pr_k[B^{e_k} \to 1] - \Pr\left[B^{0^{2^\kappa}} \to 1\right] = \frac{1}{2}\mathbb{E}_O\left[2\Delta(O)^2\right] = \mathbb{E}_O\left[\mathsf{Adv}(A, O)^2\right].$$

Since B runs A twice, by the BBBV theorem [BBBV97],

$$\mathbb{E}_O\left[\mathsf{Adv}(A, O)^2\right] = \Pr_k[B^{e_k} \to 1] - \Pr\left[B^{0^{2^\kappa}} \to 1\right] = \mathcal{O}\left(\frac{q^2}{2^\kappa}\right).$$

By Markov inequality, we have

$$\Pr_O\left[|\mathsf{Adv}(A, O)| \geq 2^{-\kappa/4}\right] = \Pr_O\left[\mathsf{Adv}(A, O)^2 \geq 2^{-\kappa/2}\right] \leq \mathcal{O}\left(\frac{q^2}{2^{\kappa/2}}\right).$$

Because q is polynomially bounded and $\kappa = \omega(\log \lambda)$, the last quantity is at most $2^{-\kappa/4}$ for all sufficiently large λ. Moreover, $\sum_\lambda 2^{-\kappa(\lambda)/4}$ converges. The Borel-Cantelli lemma ensures that, with probability 1 over O,

$$|\mathsf{Adv}(A, O)| \leq 2^{-\kappa/4}$$

holds for all but finitely many $\lambda \in \mathbb{N}$. As there are only countably many uniform adversaries A with polynomially many queries, this concludes the existence of the PRFSGs in the isometry CHFS model. The security proof for the unitary CHFS model works by replacing O with S. □

5 Oracle Separation of QPRGs from PRFSGs

This section presents a separation between QPRGs and PRFSGs with various output lengths.

Theorem 5.1. *Fix $L \leq \lambda$ such that $L = \Omega(\log \lambda)$.*[19] *There exists a unitary oracle O such that relative to $(O, \mathbf{PSPACE})$, there exist (adaptively-secure quantum-accessible) PRFSGs with output length L but QPRGs do not exist.*

In the proof of this theorem, we use a novel geometric theorem about the product Haar measure on states, which we call the barrier theorem. We first present the oracle space which we work on, then present the barrier theorem and its proof in Section 5.1. The proof of the main separation is in Section 5.2.

Oracles and PRFSGs. Fix $\ell \leq \lambda$ such that $\ell = \Omega(\log \lambda)$. We set $L = \ell(2\lambda)$ and consider the length-ℓ unitary CHFS oracle $O = O^\Phi$ for $\Phi = \{|\phi_x\rangle\}_{x \in \{0,1\}^*}$[20] sampled from a specific distribution to be specified and the **PSPACE** oracle for a fixed PSPACE-complete problem.

We consider the following PRFSGs suggested in Section 4.2 with $\kappa = m = \lambda$:

$\mathsf{Gen}^O(k, x)$**:** It outputs $|\phi_{k,x}\rangle$.

We excerpt the following lemma from the proof of Theorem 4.4 for later use. The final statement is proven using the series $\sum_{i=n}^\infty a^i = a^n/(1-a)$ for $a = 2^{-1/4}$.

[19] A standard padding argument extends this for all arbitrary polynomially bounded L with $L = \Omega(\log \lambda)$.

[20] For convenience, we omit the ancilla $|1\rangle$.

Lemma 5.2. *Fix $\kappa = m = \lambda$. For any oracle QPT adversary $\mathcal{A}$ against Gen^O, let $\mathrm{E}_{\mathcal{A}}[O, \lambda]$ be the event that*

$$\left|\Pr\left[\mathcal{A}^{\mathsf{Gen}(k,\cdot),O} \to 1\right] - \Pr_{G_{\mathsf{Haar}}}\left[\mathcal{A}^{G_{\mathsf{Haar}}(\cdot),O} \to 1\right]\right| \geq 2^{-\lambda/4}$$

happens. Then, $\Pr_{O_{2\lambda}}[\mathrm{E}_{\mathcal{A}}[O, \lambda]|O_{\neq 2\lambda}] \leq 2 \cdot 2^{-\lambda/4}$ for sufficiently large λ for any fixed $O_{\neq 2\lambda} = (O_n)_{n\neq 2\lambda}$. Furthermore, for any $n < m$ and for any fixed $O_{<n} \cup O_{>m} = (O_i)_{i<n \text{ or } i>m}$,

$$\Pr_O\left[\mathrm{E}_{\mathcal{A}}[O, i] \text{ occurs for some } n/2 \leq i < m/2 | O_{<n} \cup O_{>m}\right] \leq 11 \cdot 2^{-n/8}.$$

5.1 The barrier theorem

Let $X = \mathbb{S}(2^{n_1}) \times \cdots \times \mathbb{S}(2^{n_k})$ be the product space of quantum states (plus global phases) equipped with the product Haar measure $\sigma := \sigma_{n_1} \times \cdots \times \sigma_{n_k}$. For two elements $\Phi = (|\phi_1\rangle, \ldots, |\phi_k\rangle), \Psi = (|\psi_1\rangle, \ldots, |\psi_k\rangle)$ in X, we define the max-ℓ_2 distance $d_\infty(\Phi, \Psi) := \max_{i\in[k]} \|\phi_i - \psi_i\|_2$. For two subsets S, T of X, we define their distance as $d_\infty(S, T) := \inf_{\Phi\in S, \Psi\in T} d_\infty(\Phi, \Psi)$.

We will show the following theorem.

Theorem 5.3 (Barrier theorem)**.** *Let $X = \mathbb{S}(2^{n_1}) \times \cdots \times \mathbb{S}(2^{n_k})$ with the corresponding product Haar measure $\sigma = \sigma_{n_1} \times \cdots \times \sigma_{n_k}$, and let S_0, S_1 be two measurable subsets of X. If $d_\infty(S_0, S_1) \geq \Delta$ and $\min(\sigma(S_0), \sigma(S_1)) \geq \Gamma$, then $\sigma(X \setminus (S_0 \cup S_1)) = \Omega(\Delta\Gamma)$.*

Intuitively, Theorem 5.3 states that if two sets have a distance gap between them, then there must be a non-negligible *barrier* of the whole space that they do not cover regardless of their shape.

Proof overview. A worst-case candidate is to take a very small S_0 and the maximal S_1 subject to the conditions. Given this, a natural approach is

(i) to show that the surface area of S_0 or larger regions is bounded below, say by ε;

(ii) to "integrate" the surface area along a path from S_0 to S_1 of length Δ, obtaining a $\Omega(\varepsilon\Delta)$ lower bound.

This approach indeed works with some results from differential geometry. The most technical part is Item (i), which is formalized by the notion of Cheeger isoperimetric constant. Roughly, the Cheeger isoperimetric constant h_X ensures that:

$$\text{any "surface" area between } S_0 \text{ and } S_1 \text{ is bounded below by } (\varepsilon =) h_X \cdot \Gamma$$

given $\sigma(S_0), \sigma(S_1) \geq \Gamma$. This yields a lower bound of $h_X\Gamma\Delta$.

The proof is completed by showing the dimension-free lower bound of h_X. To do so, we use the results of Lichnerowicz and Buser, which assert that the Ricci curvature lower bound $\mathbf{Ric}_X \geq c \cdot g_X$ for $c > 0$ implies the lower bound $h_X = \Omega(\sqrt{c})$, where g_X is the Riemannian metric of the ambient space. The condition can be easily shown using some basic properties of the Ricci curvature of the product space(Eq. (9)). This differential geometric analysis shows the result for the ℓ_2 of the arc-cosine distance[21], and the desired result with the max-ℓ_2 distance follows by simple calculation.

The formal proof is given below.

Proof. Recall that $\mathbb{S}(2^n) := \{|\psi\rangle \in \mathbb{C}^{2^n} : \||\psi\rangle\|_2 = 1\}$ denotes the space of n-qubit quantum states together with the global phase, i.e., the complex hypersphere in $\mathbb{C}^{2^n}$. Our theorem is stated as follows. For two elements $\Phi = (|\phi_1\rangle, \ldots, |\phi_k\rangle), \Psi = (|\psi_1\rangle, \ldots, |\psi_k\rangle)$ in $X = \mathbb{S}(2^{n_1}) \times \cdots \times \mathbb{S}(2^{n_k})$, the max-$\ell_2$ distance is defined by $d_\infty(\Phi, \Psi) := \max_{i\in[k]} \|\phi_i - \psi_i\|_2$.

[21] which is the geodesic distance corresponding to the natural round Riemannian metric on X.

Distances. We consider various distances of quantum states. For two states in $\mathbb{S}(2^{n_i})$, the distance is

$$\|\phi - \psi\|_2 = \sqrt{2 - 2\operatorname{Re}\langle\phi|\psi\rangle}.$$

On the complex unit sphere $\mathbb{S}(2^{n_i})$, the round geodesic distance is defined by

$$d_{geo}(\phi, \psi) = \arccos(\operatorname{Re}\langle\phi|\psi\rangle).$$

We have $\|\phi - \psi\|_2 = 2\sin\left(\frac{d_{geo}(\phi,\psi)}{2}\right) \leq d_{geo}(\phi, \psi)$ in $\mathbb{S}(2^{n_i})$.

In the product space X, we consider two distances. The first distance is the max-ℓ_2 distance d_∞ defined above. The other distance is the ℓ_2-geodesic distance

$$d_2^{geo}(\Phi, \Psi) = \sqrt{\sum_{i=1}^{k} d_{geo,i}(\phi_i, \psi_i)^2} \tag{8}$$

where $d_{geo,i}$ is the angular geodesic distance in $\mathbb{S}(2^{n_i})$. We stress that the distance d_2^{geo} is defined for the geodesic distance, while d_∞ is defined for the ℓ_2 distance. We will later use the following inequality:

$$d_\infty(\Phi, \Psi) = \max_i \|\phi_i - \psi_i\|_2 \leq \max_i d_{geo,i}(\phi_i, \psi_i) \leq d_2^{geo}(\Phi, \Psi).$$

Riemannian manifolds and metrics, and Ricci curvature. A Riemannian manifold is a pair of (M, g), where M is a smooth manifold and g is a Riemannian metric[22], which assigns a positive-definite symmetric bilinear form $g_p : T_pM \times T_pM \to \mathbb{R}$ for each point $p \in M$, where T_pM denotes the tangent space of M at p. The Riemannian manifold gives rise to the geodesic distance d_g. In the complex sphere $\mathbb{S}(2^n) \subseteq \mathbb{C}^{2^n} \subseteq \mathbb{R}^{2^{n+1}}$, we use the round Riemannian metric $g_{\mathbb{S}}$ induced by the real Euclidean inner product

$$\langle u, v\rangle_{\mathbb{R}} := \operatorname{Re}\langle u|v\rangle .$$

The corresponding geodesic distance is, as discussed above,

$$d_{\mathbb{S}}(\phi, \psi) = \arccos(\operatorname{Re}\langle\phi|\psi\rangle).$$

There is a natural notion of the products of Riemannian manifolds; the geodesic distance of the products of the complex spheres with the round Riemannian metric gives the ℓ_2 distance defined in Eq. (8). Looking ahead, we will prove the statement for d_2^{geo} (instead of d_∞) using the results from differential geometry below.

On the Riemannian manifold, the Ricci curvature tensor $\mathbf{Ric}$ is uniquely determined and gives a symmetric bilinear form $\mathbf{Ric}_p : T_pM \times T_pM \to \mathbb{R}$ for each point $p \in M$. In the product of Riemannian manifolds $S = \prod_{1\leq i\leq r}(M_i, g_i)$ and the point $\mathbf{p} \in S$, the Ricci curvature on a tangent vector $\mathbf{v} = (v_1, ..., v_r) \in T_{\mathbf{p}}S$ satisfies

$$\mathbf{Ric}_S(\mathbf{v}, \mathbf{v}) = \sum_{i=1}^{r} \mathbf{Ric}_{M_i}(v_i, v_i). \tag{9}$$

The intrinsic Riemannian metric $\mathbf{g} = (g_1, ..., g_r)$[23] satisfies a similar equality $g(\mathbf{v}, \mathbf{v}) = \sum_{i=1}^{r} g_i(v_i, v_i)$.

For a constant c, we occasionally use the notation

$$\mathbf{Ric} \geq c \cdot g \iff \mathbf{Ric}_p - c \cdot g_p \succcurlyeq 0 \quad \forall p \in M \tag{10}$$

where $\succcurlyeq$ denotes the positive semi-definite inequality. Equivalently, it means

$$\mathbf{Ric}_p(v, v) \geq c \cdot g_p(v, v)$$

[22] We stress that two notions, *distance* and *metric*, are used differently; the metric is used only for the Riemannian metric.
[23] Occasionally written $\oplus_i g_i$.

holds (as an inequality over real numbers) for all $p \in M$ and $v \in T_pM$. We note that $\mathbf{Ric} \geq c \cdot g$ is usually denoted by $\mathbf{Ric} \geq c$ in the literature.

We only use Riemannian manifolds and Ricci curvatures for our underlying space $X = \mathbb{S}(2^{n_1}) \times \cdots \times \mathbb{S}(2^{n_k})$ regarding the above relations, where we recall $\mathbb{S}(2^n)$ can be understood as the real unit sphere of dimension $m_n := 2^{n+1} - 1$. The round sphere of real dimension m satisfies $\mathbf{Ric} = (m-1) \cdot g$ (i.e., an Einstein manifold) [KN96]. Therefore $\mathbf{Ric}_{\mathbb{S}(2^{n_i})} = (2^{n_i+1} - 2) \cdot g_i$.

For our main interest X together with the product round Riemannian metric, for $\mathbf{v} = (v_1, ..., v_k) \in T_{\mathbf{p}}X$ for some $\mathbf{p} \in X$, it holds that

$$\begin{aligned}
\mathbf{Ric}_X(\mathbf{v}, \mathbf{v}) &= \sum_{i=1}^{k} \mathbf{Ric}_{\mathbb{S}(2^{n_i})}(v_i, v_i) \\
&= \sum_{i=1}^{k} (2^{n_i+1} - 2) \cdot g_i(v_i, v_i) \\
&\geq 2 \sum_{i} g_i(v_i, v_i) = 2\mathbf{g}(\mathbf{v}, \mathbf{v})
\end{aligned}$$

where we use $2^{n_i+1} - 2 \geq 2$ for $n_i \geq 1$. This means X satisfies Eq. (10) for $c = 2$.

Lichnerowicz's theorem [Lic58] states that (See e.g., [Aub98, Theorem 4.19]) for an n-dimensional compact Riemannian manifold (M_n, g), if $\mathbf{Ric} \geq c \cdot g$ for some $c > 0$, then the first nonzero eigenvalue of the Laplacian[24] satisfies

$$\lambda_1 \geq \frac{nc}{n-1} \geq c. \tag{11}$$

In particular, for X together with the product round Riemannian metric, it holds that

$$\lambda_1 \geq 2. \tag{12}$$

Cheeger isoperimetric constant. We use the modern exposition for isoperimetric inequalities from [Rit23]. Consider a Riemannian manifold (M, g) with the geodesic distance d and probability measure σ (i.e., $\sigma(M) = 1$). We define the outer Minkowski boundary measure (or the *area*) of a measurable set E by

$$\sigma^+(E) = \liminf_{r \downarrow 0} \frac{\sigma(E[r]) - \sigma(E)}{r}$$

where $E[r] := \{p \in M : d(p, E) \leq r\}$ is a closed neighborhood of E. The Cheeger isoperimetric constant is

$$h_M := \inf_{E:\text{measurable}, 0<\sigma(E)<1} \frac{\sigma^+(E)}{\min(\sigma(E), 1 - \sigma(E))}. \tag{13}$$

Let λ_1 be the smallest nonzero eigenvalue of the Laplacian of M. Buser's inequality [Bus82] says that if $\mathbf{Ric} \geq -(n-1)a^2$ for some $a \geq 0$, then it holds that $\lambda_1 \leq 2a(n-1)h_M + 10h_M^2$. If $\mathbf{Ric} \geq 0$, we have $\lambda_1 \leq 10h_M^2$. This implies that, for our X,

$$h_X \geq \sqrt{1/5} \tag{14}$$

holds because $\mathbf{Ric} \geq 2\mathbf{g} \geq 0$ in X and Eq. (12).

[24] Here, the Laplacian refers the Laplace-Beltrami operator of Riemannian manifold is defined as the negative of the divergence of the gradient, i.e., $-\Delta f = -\mathsf{div}(\mathsf{grad} f)$. Negative sign is the analysts' convention which ensures the non-negativeness of the eigenvalues. The first nonzero eigenvalue is the smallest positive eigenvalue of $-\Delta$, i.e., the smallest $c > 0$ such that $-\Delta f = cf$ for some function f.

Proof of Theorem 5.3. We first work on the product space X together with the product round Riemannian metric and the corresponding ℓ_2 distance d_2. Let $v(r) := \sigma(S_0[r])$ for $0 \le r \le \Delta$ where $S_0[r] = \{x : d(x, S_0) \le r\}$. Here, we assume that v is differentiable so that $v'(r) = \sigma^+(S_0[r])$. This significantly simplifies the proof of Eq. (16), and we provide the formal proof without this assumption in Section A.1.

Note that in the definition of $S_0[r]$ in X, we use the natural distance d_2 induced by the Riemannian manifold X. Given the condition $d_\infty(S_0, S_1) \ge \Delta$, we also have $d_2^{geo}(S_0, S_1) \ge d_\infty(S_0, S_1) \ge \Delta$. For any $0 \le r < \Delta$, $v(r) \ge v(0) \ge \Gamma$. Also, $S_0[r] \cap S_1 = \emptyset$ holds so that $v(r) = \sigma(S_0[r]) \le 1 - \sigma(S_1)$ and $1 - v(r) \ge \sigma(S_1) \ge \Gamma$. This implies $\min(v(r), 1 - v(r)) \ge \Gamma$.

The Cheeger constant ensures

$$v'(r) \ge h_X \min(v(r), 1 - v(r)) \ge h_X \Gamma. \tag{15}$$

Integrating v' from 0 to $r < \Delta$ gives:

$$1 - \sigma(S_1) - \sigma(S_0) \ge \sigma(S_0[r]) - \sigma(S_0) = v(r) - v(0) \ge \int_0^r v'(x)dx \ge h_X \Gamma r. \tag{16}$$

By taking $r \to \Delta$ from below, this gives:

$$\sigma(X \setminus (S_0 \cup S_1)) = 1 - \sigma(S_0) - \sigma(S_1) \ge h_X \Gamma \Delta$$

where $h_X \ge \sqrt{1/5}$ by Eq. (14). This proves the desired result. □

5.2 Impossibility of QPRGs

We can now prove the main result of this section.

Proof of Theorem 5.1. We will construct (by diagonalization) a single unitary oracle O such that, relative to $(O, \mathbf{PSPACE})$, (i) short-PRFSGs exist for all sufficiently large λ, yet (ii) no pseudodeterministic QPRG exists (i.e., for every candidate pseudodeterministic QPRG G, there are infinitely many security parameters λ for which G is either non-pseudodeterministic or insecure). In the following, the unitary CHFS oracle S is indexed by an *oracle-input length* $t \in \mathbb{N}$: on query inputs $x \in \{0,1\}^t$ (possibly in superposition), the oracle applies a reflection $S_{|\psi_x\rangle|1\rangle}$ on $\ell(t) + 1$ qubits (cf. Definition 4.3).

Without loss of generality, we work in the uniform quantum circuit model with oracle gates for S and access to $\mathbf{PSPACE}$. Thus, the set of uniform oracle-aided QPRG candidates is countable, and we fix an enumeration over these candidates as $\{G_j\}_{j \in \mathbb{N}}$. For each G_j, let $p_j(\lambda) = \max(T_j(\lambda), \lambda)$ where $T_j(\lambda) = \mathsf{poly}(\lambda)$ is a polynomial upper bound on the running time of $G_j(1^\lambda, \cdot)$ (including oracle queries). Fix a surjection $u : \mathbb{N} \to \mathbb{N}$ such that every $j \in \mathbb{N}$ has infinitely many preimages:[25]

$$\forall j \in \mathbb{N}, \quad |\{i \in \mathbb{N} : u(i) = j\}| = \infty.$$

Intuitively, we diagonalize against the candidate $G_{u(i)}$ at security parameter λ_i.

We will define O for each index $i \in \mathbb{N}$. For index i, we choose a security parameter λ_i and then fix the oracle O on a block of oracle-input lengths $t \in I_i := [a_i, b_i]$, which are defined as follows. Choose λ_1 sufficiently large, and set $\lambda_{i+1} := 2^{2^{\lambda_i}}$ for $i \ge 1$.

Define the parameters for each i as follows:

- $\varepsilon_{u(i)}(\lambda) := \frac{d_{u(i)}}{p_{u(i)}(\lambda)^8}$ for a constant $d_{u(i)} > 0$ to be determined later.
- $c_{u(i)} \ge 1$ be a constant so that $\lambda_i^2 \cdot 2^{-c_{u(i)}/8} < \varepsilon_{u(i)}(\lambda_i)$
- $a_i := \lceil c_{u(i)} \log \lambda_i \rceil$

[25] E.g., any u such that $u(n^2 + j) = j$ for $j \le n$ works.

- $b_i := \max(p_{u(i)}(\lambda_i),\, a_i)$

Note that $b_i = \max(p_{u(i)}(\lambda_i), a_i)$ while $a_{i+1} = \Theta(\log \lambda_{i+1}) = \Theta(2^{\lambda_i})$, hence $b_i < a_{i+1}$ for all i, so the blocks I_i are disjoint and strictly increasing.

In the following, we assume that an oracle O_i is sampled so that it has the desired properties (to be explained) up to the i-th block. Our goal is to sample a new oracle that satisfies the properties up to the $(i+1)$-th block, without hurting the previous blocks.

Let $S^{(<t)}$ denote the restriction of O to oracle-input lengths $< t$, and define the conditional distribution $\mathcal{D}_{S^{(<t)}}$ in which we sample the CHFS oracles according to the product Haar random distribution conditioned on the fixed prefix $S^{(<t)}$.

The procedure is as follows. Given O_i, we fix $S^{(<a_i)}$. This will never change in the later updates. We will sample $O = O_{i+1}$ from $\mathcal{D}_{S^{(<a_i)}}$. All probabilities below are with respect to this conditional distribution.

We fix an adversary $\mathcal{A}$ for the PRFSG. Define the following event.

$A_i[\mathcal{A}, O]$**:** This is the event that $\mathcal{A}$ relative to $(O, \mathbf{PSPACE})$ against Gen^O has an advantage at most $2^{-\lambda/4}$ for the security game of PRFSGs with input length λ for all $2\lambda \in [a_i, a_{i+1} - 1] (\supset I_i)$.

Note that this event does not occur whenever $\mathrm{E}_{\mathcal{A}}[O, \lambda]$ occurs for some $2\lambda \in [a_i, a_{i+1} - 1]$. By Lemma 5.2, we have:

$$\Pr[\neg A_i[\mathcal{A}, O]] \le 11 \cdot 2^{-a_i/8} \le 11\lambda_i^{-c_{u(i)}/8} < \frac{\varepsilon_{u(i)}(\lambda_i)}{\lambda_i^2}, \tag{17}$$

for sufficiently large i, where we use $a_i = \lceil c_{u(i)} \log \lambda_i \rceil$ and the choice of $c_{u(i)}$.

Now we analyze the candidate $G_{u(i)}$ relative to $O \leftarrow \mathcal{D}_{S^{(<a_i)}}$. We consider its security at the security parameter λ_i. Define the following event.

$B_i[O]$**:** This is the event that $G_{u(i)}$ is pseudodeterministic with probability at least $1 - \varepsilon_{u(i)}(\lambda_i)$ over the random choice of O (as in Definition 3.8) at parameter λ_i. In other words, this event occurs only when: for at least a $(1 - \varepsilon_{u(i)}(\lambda_i))$-fraction of $x \in \{0,1\}^{n(\lambda_i)}$, there exists $y_{x,O}$ such that $\Pr\left[G^O(x) \to y_{x,O}\right] \ge 1 - \varepsilon_{u(i)}(\lambda_i)$.

Note that $\neg B_i[O]$ means that $G_{u(i)}$ fails to achieve the $(1 - \varepsilon_{u(i)})$-pseudodeterminism at least for the parameter λ_i. We will need the following lemma whose proof is deferred to Section 5.3.

Lemma 5.4. *Suppose that O is sampled from $\mathcal{D}_{S^{(<a_i)}}$. Let $G = (G_\lambda)_\lambda$ be a candidate oracle QPRG that takes an $n = n(\lambda)$-bit seed and outputs an $m = m(\lambda)$-bit string relative to $(O, \mathbf{PSPACE})$ that makes at most $p = \mathsf{poly}(\lambda)$ queries.*

There exists an oracle QPT algorithm $\mathcal{B}_G$ having access to **PSPACE** *against the QPRG security such that: If G is $(1 - \varepsilon(\lambda))$-pseudodeterministic with probability at least $(1 - \varepsilon(\lambda))$ over O for parameter $\lambda = \lambda_i$, then, it holds that*

$$\Pr_O\left[\mathcal{B}_G^{\mathbf{PSPACE}} \text{ against } G^O \text{ has advantage } 0.1 \text{ at parameter } \lambda_i\right] \ge 1 - \frac{1}{\lambda_i^2}.$$

Roughly, it says that there is an adversary breaking the pseudorandomness of $G_{u(i)}$ with very high probability over the oracle, as long as $\Pr[B_i[O]]$ is big enough.

Regarding this, we consider the following two cases. We choose the oracle in the next block depending on them.

Case 1: $\Pr[\neg B_i[O]] \ge \varepsilon_{u(i)}(\lambda_i)$. In this case, we sample the oracle O from $\mathcal{D}_{S^{(<a_i)}}$ conditioned on $\neg B_i[O]$. This condition enforces that, for any $\mathcal{A}$,

$$\Pr[A_i[\mathcal{A}, O] | \neg B_i[O]] \ge 1 - \frac{1}{\lambda_i^2} \ge 1 - \frac{1}{i^2} \tag{18}$$

holds for sufficiently large i.

Case 2: $\Pr[\neg B_i[O]] < \varepsilon_{u(i)}(\lambda_i)$. In this case, we sample the oracle O from $\mathcal{D}_{S^{(<a_i)}}$ (without any condition). In this case, Lemma 5.4 says that the distinguisher $\mathcal{B} = \mathcal{B}_G$ achieves an advantage 0.1 at parameter λ_i with probability at least $1 - \frac{1}{\lambda_i^2} \geq 1 - \frac{1}{i^2}$ for sufficiently large i. Also, in this case, it holds that

$$\Pr[A_i[\mathcal{A}, O]] \geq 1 - \frac{\varepsilon_{u(i)}(\lambda_i)}{\lambda_i^2} \geq 1 - \frac{1}{i^2}. \tag{19}$$

In any case, we define $S^{(<a_{i+1})}$ by the chosen O up to input length less than a_{i+1}. Our final oracle will be the oracle matching to $S^{<a_i}$ up to length less than a_i for all $i \in \mathbb{N}$, which is iteratively constructed by the above procedure.

The following claim holds because, for any $\mathcal{A}$, regarding Eqs. (18) and (19) with $\sum_i 1/i^2 < \infty$, the Borel-Cantelli lemma says that the event $\neg A_i[\mathcal{A}, O]$ occurs only finitely many times.

Claim 5.5. *Relative to O sampled above and* **PSPACE**, Gen^O *is a secure PRFSG with probability 1.*

Fix a QPRG candidate $G = G_i$, and write $u^{-1}(i) = \{i_1, i_2, ...\}$. Because $\neg B_{i_j}[O]$ means G_i does not satisfy $(1 - \varepsilon_{u(i)})$-pseudodeterminism, the following is clear.

Claim 5.6. *Relative to O sampled above and* **PSPACE**, *if* **Case 1** *occurs infinitely many times in $u^{-1}(i)$, G does not satisfy the $(1 - \varepsilon_{u(i)})$-pseudodeterminism. In particular, it cannot be a QPRG with negligible errors.*

On the other hand, Lemma 5.4 shows the following.

Claim 5.7. *Relative to O sampled above and* **PSPACE**, *if* **Case 2** *occurs infinitely many times in $u^{-1}(i)$, $\mathcal{B}_G$ breaks the pseudorandomness of G.*

This completes the proof. □

5.3 Proof of Lemma 5.4

We will first state and prove the following lemma, which will be used in the proof of Lemma 5.4. It essentially says that if a quantum algorithm outputs a bit b_Φ with high probability over the choice of the CHFS oracle Φ, then this output bit is independent of the oracle with high probability over Φ.

As in the setting of Lemma 5.4, we fix the CHFS oracle O up to the length $< a_i$. For convenience, we write $\Phi \leftarrow \sigma$ to denote $S = S^\Phi$ is sampled from the conditional distributions $\mathcal{D}_{S^{(<a_i)}}$.

Lemma 5.8. *Let $\mathcal{S}$ be the (unitarized) quantum-accessible CHFS oracle with $\ell(\lambda) \leq \lambda$ and let $A^{\mathcal{S}}$ be a polynomial-query oracle algorithm making at most $T = T(\lambda)$ oracle queries. Let $p = \mathsf{poly}(\lambda)$ be the maximal length of the CHFS oracles that A accesses.*[26] *Suppose that there exist a function $\varepsilon = \varepsilon(\lambda)$ and a bit $b_\Phi \in \{0,1\}$ such that*

$$\Pr_{\Phi \leftarrow \sigma}\left[\Pr\left[A^{S^\Phi}(1^\lambda) \to b_\Phi\right] \geq \frac{2}{3}\right] = 1 - \varepsilon. \tag{20}$$

Then there exists $b \in \{0,1\}$[27] *such that $\Pr_{\Phi \leftarrow \sigma}[b = b_\Phi] = 1 - \mathcal{O}(T\varepsilon)$.*

In particular, if $\varepsilon(\lambda) = \mathsf{negl}(\lambda)$ then $\Pr[b = b_\Phi] = 1 - \mathsf{negl}(\lambda)$, and if $\varepsilon(\lambda) = 1/\mathsf{poly}(\lambda)$ then $\Pr[b = b_\Phi] = 1 - 1/\mathsf{poly}(\lambda)$.

Proof. Given the running time bound p, the algorithm only accesses the oracle up to the maximum query length p. Let $X = \mathbb{S}(2^{n_1}) \times \cdots \times \mathbb{S}(2^{n_k})$ be the states[28] to define the CHFS oracle up to the length p, with the corresponding product Haar measure $\sigma = \sigma_{n_1} \times \cdots \times \sigma_{n_k}$. Let $S_0, S_1 \subseteq X$ be defined as

$$S_0 := \left\{\Phi \in X : \Pr\left(A^{S^\Phi}(1^\lambda) \to 0\right) \geq 2/3\right\}, \quad S_1 := \left\{\Phi \in X : \Pr\left(A^{S^\Phi}(1^\lambda) \to 1\right) \geq 2/3\right\}.$$

[26] This premise holds for any polynomial space algorithms.
[27] The bit b might depend on the parameter λ, but we suppress this dependency for simplicity.
[28] This is implicitly parameterized by λ.

By the hypothesis in Eq. (20), with probability at least $1-\varepsilon(\lambda)$ over $\Phi \leftarrow \sigma$, either

$$\Pr\Big[A^{S^\Phi} \to 1\Big] \geq 2/3 \quad \text{or} \quad \Pr\Big[A^{S^\Phi} \to 0\Big] \geq 2/3,$$

thus

$$\sigma(X \setminus (S_0 \cup S_1)) \leq \varepsilon(\lambda). \tag{21}$$

Choose arbitrary $\Phi \in S_0$ and $\Psi \in S_1$. The diamond distance between $\mathcal{S}^\Phi$ and $\mathcal{S}^\Psi$ can be bounded by

$$\begin{aligned}
\|S^\Phi(\cdot) - S^\Psi(\cdot)\|_\diamond &\leq 2\|S^\Phi - S^\Psi\|_{op} \\
&= 2\|\sum_x |x\rangle\langle x| \otimes (S_{|\phi_x\rangle} - S_{|\psi_x\rangle})\|_{op} \\
&= 2\max_x \|(S_{|\phi_x\rangle} - S_{|\psi_x\rangle})\|_{op} \\
&= 4\max_x \| \, |\phi_x\rangle\langle\phi_x| - |\psi_x\rangle\langle\psi_x| \, \|_{op} \\
&\leq 4\max_x \|\phi_x - \psi_x\|_2 = 4d_\infty(\Phi,\Psi)
\end{aligned}$$

where we use Eqs. (5) and (6) and

$$\| \, |\phi_x\rangle\langle\phi_x| - |\psi_x\rangle\langle\psi_x| \, \|_{op} \leq \| \, |\phi_x\rangle - |\psi_x\rangle \, \|_2 \leq \|\phi_x - \psi_x\|_2.$$

This implies that the diamond distance of the unitary oracles S^Φ and S^Ψ is bounded by $\|S^\Phi(\cdot) - S^\Psi(\cdot)\|_\diamond \leq 4d_\infty(\Phi,\Psi)$. On one hand, if the algorithm A makes T queries to the oracles, the subadditivity of the diamond norm under composition implies that

$$\|A^{S^\Phi} - A^{S^\Psi}\|_\diamond \leq 4T \cdot d_\infty(\Phi,\Psi).$$

On the other hand, by the definition of the diamond distance, we can lower bound this quantity by

$$\|A^{S^\Phi} - A^{S^\Psi}\|_\diamond \geq \Big|\Pr\Big(A^{S^\Phi} \to 1\Big) - \Pr\Big(A^{S^\Psi} \to 1\Big)\Big| \geq \frac{1}{3},$$

where the last inequality is obtained from $\Phi \in S_0$ and $\Psi \in S_1$. Given it holds for arbitrary $\Phi \in S_0$ and $\Psi \in S_1$, this implies that

$$\Delta = d_\infty(S_0, S_1) \geq \frac{1}{12T} \tag{22}$$

which is inverse-polynomially bounded. Let $\Gamma := \min(\sigma(S_0), \sigma(S_1))$. We find therefore ourselves in the hypothesis of Theorem 5.3, thus there exists an absolute constant $\kappa > 0$ such that

$$\sigma\big(X \setminus (S_0 \cup S_1)\big) \geq \kappa \cdot \Gamma \cdot \Delta. \tag{23}$$

We upper bound Γ by contradiction. Fix a sufficiently large absolute constant $K > 0$ and suppose that

$$\Gamma \;>\; K\,T\,\varepsilon. \tag{24}$$

Combining (24) with (22) and (23) gives

$$\sigma\big(X \setminus (S_0 \cup S_1)\big) \;>\; \kappa \cdot (K\,T\,\varepsilon) \cdot \frac{1}{12T} \;=\; \frac{\kappa K}{12}\,\varepsilon.$$

Choosing K large enough so that $\frac{\kappa K}{12} > 1$ yields $\sigma(X \setminus (S_0 \cup S_1)) > \varepsilon$, contradicting (21). Hence (24) is false, i.e.,

$$\Gamma \;=\; \min(\sigma(S_0), \sigma(S_1)) \;\leq\; \mathcal{O}(T\varepsilon). \tag{25}$$

Define $b(\lambda) \in \arg\max_{b\in\{0,1\}} \Pr_{\Phi\leftarrow\sigma}[b = b_\Phi]$. Since S_0 and S_1 are disjoint and $\sigma(S_0 \cup S_1) \geq 1-\varepsilon$ by (21), we have

$$\begin{aligned}\Pr_{\Phi\leftarrow\sigma}[\,b = b_\Phi\,] \geq \sigma(S_b) &= \sigma(S_0 \cup S_1) - \sigma(S_{1-b}) \\ &\geq (1-\varepsilon) - \min(\sigma(S_0), \sigma(S_1)) \\ &\geq 1 - \varepsilon - \mathcal{O}(T\varepsilon),\end{aligned}$$

where the last step uses (25). This concludes the proof. □

We also need the following lemma on t-designs, which will be used in the proof of Lemma 5.4 to simulate the CHFS oracle.

Definition 5.9 (Approximate design [AE07])**.** *A probability distribution S over $\mathbb{S}(N)$ is an* ε-approximate t-design *state if:*

$$(1-\varepsilon)\mathbb{E}_{|\psi\rangle\sim\sigma_N}\left[(|\psi\rangle\langle\psi|)^{\otimes t}\right] \preceq \mathbb{E}_{|\psi\rangle\sim S}\left[(|\psi\rangle\langle\psi|)^{\otimes t}\right] \preceq (1+\varepsilon)\mathbb{E}_{|\psi\rangle\sim\sigma_N}\left[(|\psi\rangle\langle\psi|)^{\otimes t}\right].$$

Definition 5.10 (Phase-invariant design)**.** *An ε-approximate t-design state S is called a phase-invariant design if: For all $0 \leq i, j \leq t$ and $i \neq j$, with convention $|\psi\rangle^0 = |0\rangle$ and $\langle\psi|^0 = \langle 0|$,*

$$\mathbb{E}_{|\psi\rangle\sim S}\left[|\psi\rangle^{\otimes i}\langle\psi|^{\otimes j}\right] = 0.$$

The following lemma is shown by the corresponding results of [Kre21] using the phase-invariant unitary designs instead of the standard designs.

Lemma 5.11 ([Kre21], adapted)**.** *For each $n, t \in \mathbb{N}$ and $\varepsilon > 0$, there exists $m \leq \mathsf{poly}\big(n, t, \log\frac{1}{\varepsilon}\big)$ and a $\mathsf{poly}\big(n, t, \log\frac{1}{\varepsilon}\big)$-time classical algorithm $\mathcal{T}$ that takes as input a random string $x \sim \{0,1\}^m$ and outputs a description of a quantum state on n qubits such that the states sampled from $\mathcal{T}$ form a phase-invariant ε-approximate t-design over $\mathbb{S}(2^n)$.*

We additionally need the following CHFS oracle simulation lemma using phase-invariant designs.

Lemma 5.12 (CHFS design simulation)**.** *Let $\mathcal{D}$ be the distribution where the states $\{|\psi_x\rangle\}_{x\in\{0,1\}^*}$ are independent phase-invariant δ-approximate $2q$-design states, and let S^Ψ be the corresponding CHFS oracle. Then every quantum algorithm A making at most q coherent queries up to input length q satisfies*

$$\left|\Pr_{\Phi\leftarrow\mathcal{S}}\left[A^{S^\Phi} \to 1\right] - \Pr_{\Psi\leftarrow\mathcal{D}}\left[A^{S^\Psi} \to 1\right]\right| \leq O\left(q \cdot C^{q^2} \cdot \delta\right)$$

for some universal constant $C > 0$ (not depending on A), where $\mathcal{S}$ is the uniform CHFS oracle distribution. The same conclusion holds if the design seeds are only $2q$-wise independent across labels.

Proof of Lemma 5.12. Consider a D-dimensional unitary U and let

$$|J(U)\rangle = (I \otimes U)|\Omega\rangle$$

be its normalized Choi state where $|\Omega\rangle = \frac{1}{\sqrt{D}}\sum_i |ii\rangle$. By [Kre21, Lemmas 23, 24], for every q-query oracle algorithm A^U, there exists some algorithm B such that

$$\Pr\left[A^U \to 1\right] = D^{2q}\Pr\left[B(|J(U)\rangle^{\otimes q}) \to 1\right].$$

This gives, by the monotonicity of the trace distance, we have

$$\left|\Pr_{\Phi\leftarrow\mathcal{S}}\left[A^{S^\Phi} \to 1\right] - \Pr_{\Psi\leftarrow\mathcal{D}}\left[A^{S^\Psi} \to 1\right]\right|$$

$$= O\left(D^{2q} \cdot \left\| \mathbb{E}_{\Phi \leftarrow \mathcal{S}} \left|J(S^\Phi)\right\rangle\!\left\langle J(S^\Phi)\right|^{\otimes q} - \mathbb{E}_{\Psi \leftarrow \mathcal{D}} \left|J(S^\Psi)\right\rangle\!\left\langle J(S^\Psi)\right|^{\otimes q} \right\|_1\right),$$

where S^Φ and S^Ψ are concatenated up to the input length q, which gives $D = 2^{O(q)}$. It remains to bound the Choi-moment difference.

Note that we have less than $2^{q+1} =: M$ choices of $x \in \{0,1\}^{\leq q}$. For each x,

$$S_{\phi_x} = I - |0\rangle\!\langle 0| - |\phi_x\rangle\!\langle \phi_x| + |\phi_x\rangle\!\langle 0| + |0\rangle\!\langle \phi_x|.$$

Write $T_{x,j}$ to denote the j-th term in the above equation for $j \in [5]$, i.e., $T_{x,1} = I, T_{x,2} = |0\rangle\!\langle 0|$, etc. Let $\mathcal{I}$ be the set of all pairs of (x,j), so that $|\mathcal{I}| \leq 5M$. Expanding $\left|J(S^\Phi)\right\rangle^{\otimes q}$ can be written as the linear summation of

$$\bigotimes_{i=1}^{q} (I \otimes |x_i\rangle\!\langle x_i| \otimes T_{x_i,j_i}) \left|\Omega\right\rangle^{\otimes q}$$

of (less than) $(5M)^q$ terms. A similar result holds when expanding $\left\langle J(S^\Phi)\right|^{\otimes q}$, resulting $(5M)^{2q}$ terms total.

Fix one expanded term. It involves at most $2q$ label occurrences. For each label x, the contribution of it is a moment of the form

$$\mathbb{E}\left[|\phi_x\rangle^{\otimes a_x} \langle \phi_x|^{\otimes b_x}\right], \qquad a_x, b_x \leq 2q.$$

If $a_x \neq b_x$, both the Haar moment and the design moment vanish by the phase-invariant property. If $a_x = b_x = s$, then $s \leq 2q$, and the δ-approximate $2q$-design gives

$$\left\| \mathbb{E}_{\psi_x} |\psi_x\rangle\!\langle \psi_x|^{\otimes s} - \mathbb{E}_{\phi_x} |\phi_x\rangle\!\langle \phi_x|^{\otimes s} \right\|_1 \leq 2\delta.$$

Since the fixed expanded term contains at most $2q$ relevant labels, a simple hybrid shows that the distance is at most $4q\delta$.

By linearity of expectation, the difference of the two Choi moments is the sum, over all expanded monomials, of the corresponding differences of monomial expectations. Let $\mathsf{J}(A) := (I \otimes A)|\Omega\rangle$. For $\alpha, \beta \in \mathcal{I}^q$, define

$$T_{\alpha,\beta}(\Phi) := \left(\bigotimes_{i=1}^{q} \mathsf{J}(B_{\alpha_i}(\Phi))\right) \left(\bigotimes_{i=1}^{q} \mathsf{J}(B_{\beta_i}(\Phi))\right)^\dagger.$$

For each fixed expanded monomial $T_{\alpha,\beta}$, the previous argument shows

$$\left\| \mathbb{E}_{\Phi \leftarrow \mathcal{S}} T_{\alpha,\beta}(\Phi) - \mathbb{E}_{\Psi \leftarrow \mathcal{D}} T_{\alpha,\beta}(\Psi) \right\|_1 \leq 4q\delta.$$

There are at most $(5M)^{2q}$ such monomials. Hence, by the triangle inequality for the trace norm,

$$\begin{aligned}
&\left\| \mathbb{E}_{\Phi \leftarrow \mathcal{S}} \left|J(S^\Phi)\right\rangle\!\left\langle J(S^\Phi)\right|^{\otimes q} - \mathbb{E}_{\Psi \leftarrow \mathcal{D}} \left|J(S^\Psi)\right\rangle\!\left\langle J(S^\Psi)\right|^{\otimes q} \right\|_1 \\
&\leq \sum_{\alpha,\beta} \left\| \mathbb{E}_{\Phi \leftarrow \mathcal{S}} T_{\alpha,\beta}(\Phi) - \mathbb{E}_{\Psi \leftarrow \mathcal{D}} T_{\alpha,\beta}(\Psi) \right\|_1 \\
&\leq (5M)^{2q} \cdot 4q\delta,
\end{aligned}$$

thus the overall bound becomes

$$O\left(4q\delta \cdot (5DM)^{2q}\right) = O\left(q \cdot C^{q^2} \cdot \delta\right)$$

for some $C > 0$, because $D = 2^{O(q)}$ and $M \leq 2^{q+1}$. This proves the main upper bound. The final statement is clear because the proof only ever uses independence among the at most $2q$ labels appearing in a fixed expanded term. □

Proof of Lemma 5.4. Recall Definition 3.8 for the definition of the pseudodeterministic QPRG. Consider a QPRG candidate $G = (G_\lambda)_\lambda$ that takes an n-bit seed and outputs an m-bit string relative to the CHFS oracles and **PSPACE** that runs in polynomial queries. Let $p(\lambda)$ be the upper bound of the running time of QPT G_λ.

We say that an input $x \in \{0,1\}^{n(\lambda_i)}$ is $\varepsilon_\star$-*good* if

$$\Pr_S\left[\exists\, y_{S,x} \text{ s.t. } \Pr\left[G^{S,\mathbf{PSPACE}}(x) \to y_{S,x}\right] \geq 1-\varepsilon_\star\right] \geq 1-\varepsilon_\star. \tag{26}$$

Furthermore, if the inner event holds for a particular S (i.e., there exists such a $y_{S,x}$), we say (S,x) is *good*.

We write $I_{S,x}$ to denote the indicator that (S,x) is good. By the hypothesis that G satisfies $(1-\varepsilon(\lambda_i))$-pseudodeterminism for $(1-\varepsilon(\lambda_i))$-fraction of O, $\Pr_x[I_{S,x}=1] \geq 1-\varepsilon(\lambda_i)$, and therefore

$$\mathbb{E}_{S,x}[I_{S,x}] \geq (1-\varepsilon(\lambda_i))^2 \geq 1-2\varepsilon(\lambda_i).$$

Fix any $\varepsilon_\star = (2\varepsilon)^{0.5}$. By Markov's inequality (or the standard averaging argument),

$$\Pr_x\left[\mathbb{E}_S[I_{S,x}] \leq 1-\varepsilon_\star\right] = \Pr_x\left[1-\mathbb{E}_S[I_{S,x}] \geq \varepsilon_\star\right] \leq \frac{1-\mathbb{E}_{S,x}[I_{S,x}]}{\varepsilon_\star} \leq \frac{2\varepsilon}{\varepsilon_\star} = \varepsilon_\star.$$

Hence, for at least a $1-\varepsilon_\star$ fraction of x we have $\mathbb{E}_S[I_{S,x}] \geq 1-\varepsilon_\star$, i.e., x is $\varepsilon_\star$-good. We call such x *semi-deterministic*.

We will construct $F^{\mathbf{PSPACE}}$ that agrees with $G^{O,\mathbf{PSPACE}}$ with high probability on all semi-deterministic x. We divide O into the input length $< a_i$ (which is already fixed by $S^{(<a_i)}$) and $\geq a_i$.

We first show that $S^{(<a_i)}$ can be simulated up to negligible error in $\mathsf{poly}(\lambda_i)$ time using tomography. Recall that the unitary CHFS oracle is indexed by an oracle-input length $t < a_i$, and that on input $x \in \{0,1\}^t$, it applies a reflection $S_{t,x} = S_{|\psi_x\rangle|1\rangle}$ on $\ell(t)+1$ qubits. That is, $S_{t,x}$ acts on dimension $d(t) = 2^{\ell(t)+1} \leq 2^{t+1} \leq 2^{a_i} = \mathsf{poly}(\lambda_i)$ because $a_i = O(\log \lambda_i)$.[29]

For each fixed pair (t,x) with $t < a_i$, applying Theorem 3.19to the black-box unitary $S_{t,x}$ with $\varepsilon_{\mathsf{tom}} := \lambda_i^{-c'}$ (for a constant $c' > 0$ to be determined later) and $\delta_{\mathsf{tom}} := 2^{-2\lambda_i}$ yields an algorithm that outputs a classical description of a unitary $\widetilde{S}_{t,x}$ such that, with probability at least $1-\delta_{\mathsf{tom}}$,

$$\left|S_{t,x} - \widetilde{S}_{t,x}\right|_\diamond \leq \varepsilon,$$

using $\mathcal{O}\left(\frac{d(t)^2}{\varepsilon_{\mathsf{tom}}} \log \frac{1}{\delta_{\mathsf{tom}}}\right) = \mathsf{poly}(\lambda_i)$ queries and time $\mathsf{poly}\left(d(t), \frac{1}{\varepsilon_{\mathsf{tom}}}, \log\left(\frac{1}{\delta_{\mathsf{tom}}}\right)\right) = \mathsf{poly}(\lambda_i)$.

The number of such pairs (t,x) is $N_i = \sum_{t=1}^{a_i-1} 2^t \leq 2^{a_i} = \mathsf{poly}(\lambda_i)$. By a union bound over all N_i invocations, with probability at least $1 - N_i\delta_{\mathsf{tom}} \geq 1 - \mathsf{poly}(\lambda_i)\cdot 2^{-2\lambda_i} = 1-\mathsf{negl}(\lambda_i)$, all tomography subroutines succeed simultaneously, producing descriptions $\{\widetilde{S}_{t,x}\}_{t<a_i, x\in\{0,1\}^t}$ with diamond-error at most $\varepsilon_{\mathsf{tom}}$ each. This shows that the trace distance between the output distribution can be bounded by

$$\|G^S(\cdot) - G^{\widetilde{S}}(\cdot)\|_1 \leq p(\lambda_i)\cdot\varepsilon_{\mathsf{tom}} = p(\lambda_i)\cdot\lambda_i^{-c'} \tag{27}$$

for any input, where we omit the other oracle queries that are identical for the left and right cases. We choose c' so that $\lambda_i^{-c'} \leq \frac{1}{100p(\lambda_i)^2\lambda_i^2}$.

Now consider the oracle query length $\geq a_i$. Let $\delta = \frac{1}{100D\lambda_i^2 p(\lambda_i)^2 C^{p(\lambda_i)^2}}$ where C and D are constants from Lemma 5.12. Let $\mathcal{T}(\cdot)$ be the algorithm from Lemma 5.11 that, on input a uniform seed $r \in \{0,1\}^{s(q)}$, outputs a (description of a) phase-invariant q-qubit state whose output state forms a δ-approximate $2p(\lambda_i)$-design.[30]

For each $t \in [p(\lambda_i)]$, we sample a function $h_t : \{0,1\}^t \to \{0,1\}^{s(\ell(t))}$ from a $2p(\lambda_i)$-wise independent family (independently for different t). By [Zha12], we can assume that h_t is sampled from a truly random function as F and G make at most $p(\lambda_i)$ quantum queries.

We now describe how F simulates G's access to the unitary CHFS oracle S, given the above data. On an oracle query with input $x \in \{0,1\}^t$ (possibly in superposition), F proceeds as follows.

[29] We use $\ell(\lambda) \leq \lambda$ here. It can be relaxed to $\ell = O(\lambda)$.

[30] The seed length $s(q)$ depends also on $p(\lambda_i)$ and $\log(1/\delta)$; we suppress these dependencies for simplicity.

1. It uses $\widetilde{S}$ instead of S for the input length $t < a_i$. For the input length $t \geq a_i$, it proceeds to the next step.

2. Coherently compute $r_x := h_t(x) \in \{0,1\}^{s(\ell(t))}$.

3. Coherently compute $|\psi_x\rangle := T(r_x)$. Note that the parameters are such that on input $x \in \{0,1\}^t$, $h_t(x)$ is of size $s(\ell(t))$, which is the size of the randomness needed to generate a $\ell(t)$-qubit δ-error $2p(\lambda_i)$-design using $\mathcal{T}(\cdot)$.

4. Apply the reflection $S_{|\psi_x\rangle|1\rangle}$ on the $\ell(t)+1$-qubit register $\mathbf{Y}$ (cf. Definition 4.3).

Equivalently, instead of querying the unitary CHFS oracle, F uses

$$O' := \sum_{x\in\{0,1\}^t} |x\rangle\langle x|_{\mathbf{X}} \otimes S_{|\psi_x\rangle|1\rangle},$$

on $\rho_{\mathbf{XYZ}}$ for $|x| \geq a_i$, which is indistinguishable from the CHFS oracle up to some errors that we will now analyze.

We write $f_z(x)$ for the z-th bit of $F^{\mathbf{PSPACE}}(x)$, $g_z(x)$ for the z-th bit of $G^{S,\mathbf{PSPACE}}(x)$.

Fix a semi-deterministic x. Applying Lemma 5.8 to $g_z(x)$ (which makes at most $T \leq p(\lambda)$ oracle queries) given the fact that $G(x)$ is pseudodeterministic with probability at least $(1-\varepsilon_\star)$ over the oracles, yields that $g_z(x)$ must be fixed with probability at least $(1 - O(p(\lambda_i)\varepsilon_\star))$ over the CHFS oracles.

Combining this with the δ-approximate $2p(\lambda_i)$-design simulation error gives

$$\begin{aligned}\Pr\left[g_z(x) = f_z(x)\right] \geq\ & 1 - \mathcal{O}(p(\lambda_i)\varepsilon_\star) - O(p(\lambda_i)\cdot C^{p(\lambda_i)^2}\cdot\delta) - \frac{1}{100\lambda_i^2 p(\lambda_i)} \\ \geq\ & 1 - \mathcal{O}(p(\lambda_i)\varepsilon_\star) - \frac{1}{100\lambda_i^2 p(\lambda_i)} - \frac{1}{100\lambda_i^2 p(\lambda_i)},\end{aligned}$$

where we use Lemma 5.12 and the term $\frac{1}{100\lambda_i^2 p(\lambda_i)}$ is from the tomography error in Eq. (27). By a union bound over $z \in [m]$ and using $m \leq p(\lambda)$, we obtain

$$\Pr\left[G^{S,\mathbf{PSPACE}}(x) = F^{\mathbf{PSPACE}}(x)\right] \ \geq\ 1 - O(p(\lambda)^2\varepsilon_\star) - \frac{1}{50\lambda_i^2}, \tag{28}$$

for all $\varepsilon_\star$-semi-deterministic x.

Recall the condition $\varepsilon_\star(\lambda) = (2\varepsilon(\lambda))^{0.5} \leq \frac{\sqrt{2c}}{p(\lambda)^4}$ for some small $c > 0$. We choose c so that the term $O(p(\lambda)^2\varepsilon_\star)$ is bounded above by $\frac{1}{50\lambda_i^2}$. This shows that the right-hand side of Eq. (28) is at least $1 - \frac{1}{4\lambda_i^2}$ for all sufficiently large λ.

Now, fix the initialization randomness r of F for the tomography and designs. Below, we assume that F uses the randomness r, and consider the language

$$L_{F,r} := \left\{ y \in \{0,1\}^{m(\lambda)} : \exists x \in \{0,1\}^{n(\lambda)} \text{ s.t. } \Pr\left[F^{\mathbf{PSPACE}}(x) = y\right] > 1 - \frac{1}{4\lambda^2} \right\}.$$

For fixed (x,y), a **PSPACE** machine can approximate $\Pr\left[F^{\mathbf{PSPACE}}(x) = y\right]$ to additive error $o(\lambda^{-2})$ (since $\mathbf{BQP}^{\mathbf{PSPACE}} \subseteq \mathbf{PSPACE}$), and hence decide whether it exceeds $1 - 1/4\lambda^2$. It follows that $L \in \mathbf{PSPACE}$: on input y, enumerate all $x \in \{0,1\}^n$ using polynomial space and accept iff any x satisfies $\Pr\left[F^{\mathbf{PSPACE}}(x) = y\right] > 1 - 1/4\lambda^2$. Therefore, there exists a $\mathbf{PSPACE}$-oracle algorithm $\mathcal{B}_0^{\mathbf{PSPACE}}$ that can decide membership in L.

We define our adversary $\mathcal{B}$ formally: it first computes the tomography for the input length $< a_i$, $2p(\lambda_i)$-wise independent functions, and state designs described above as initialization. Then, given this, it runs $\mathcal{B}_0^{\mathbf{PSPACE}}$.

We consider the advantage of $\mathcal{B}$ against the pseudorandomness of G. By Eq. (28), for every semi-deterministic x it holds that

$$\Pr_S\left[\mathcal{B}^{\mathbf{PSPACE}}(G^{S,\mathbf{PSPACE}}(x))=1\right]\geq 1-\frac{1}{4\lambda_i^2}.$$

As a $(1-O(1/\lambda_i^2))$-fraction of x, with a sufficiently small constant, satisfies the above, we have

$$\Pr_{S,x}\left[\mathcal{B}^{\mathbf{PSPACE}}(G^{S,\mathbf{PSPACE}}(x))=1\right]\geq 1-\frac{1}{3\lambda_i^2}.$$

The standard averaging argument says that

$$\Pr_S\left[\left[\mathcal{B}^{\mathbf{PSPACE}}(G^{S,\mathbf{PSPACE}}(x))=1\right]\geq\frac{2}{3}\right]\geq 1-\frac{1}{\lambda_i^2}.$$

On the other hand, $|L_F|\leq 2^n$ (each x can contribute at most one such y), so for uniform $y\leftarrow\{0,1\}^m$ we have $\Pr\left[\mathcal{B}^{\mathbf{PSPACE}}(y)=1\right]\leq 2^{n-m}\leq 1/2$ (using that $m>n$ for a QPRG). This gives a constant distinguishing advantage of $\mathcal{B}$ against G for the parameter λ_i. □

Proof of Lemma 5.12. Consider a D-dimensional unitary U and let

$$|J(U)\rangle=(I\otimes U)|\Omega\rangle$$

be its normalized Choi state where $|\Omega\rangle=\frac{1}{\sqrt{D}}\sum_i|ii\rangle$. By [Kre21, Lemmas 23, 24], for every q-query oracle algorithm A^U, there exists some algorithm B such that

$$\Pr\left[A^U\to 1\right]=D^{2q}\Pr\left[B(|J(U)\rangle^{\otimes q})\to 1\right].$$

This gives, by the monotonicity of the trace distance, we have

$$\begin{aligned}&\left|\Pr_{\Phi\leftarrow\mathcal{S}}\left[A^{S^\Phi}\to 1\right]-\Pr_{\Psi\leftarrow\mathcal{D}}\left[A^{S^\Psi}\to 1\right]\right|\\&=O\left(D^{2q}\cdot\left\|\mathbb{E}_{\Phi\leftarrow\mathcal{S}}\left|J(S^\Phi)\right\rangle\!\left\langle J(S^\Phi)\right|^{\otimes q}-\mathbb{E}_{\Psi\leftarrow\mathcal{D}}\left|J(S^\Psi)\right\rangle\!\left\langle J(S^\Psi)\right|^{\otimes q}\right\|_1\right),\end{aligned}$$

where S^Φ and S^Ψ are concatenated up to the input length q, which gives $D=2^{O(q)}$. It remains to bound the Choi-moment difference.

Note that we have less than $2^{q+1}=:M$ choices of $x\in\{0,1\}^{\leq q}$. For each x,

$$S_{\phi_x}=I-|0\rangle\!\langle 0|-|\phi_x\rangle\!\langle\phi_x|+|\phi_x\rangle\!\langle 0|+|0\rangle\!\langle\phi_x|.$$

Write $T_{x,j}$ to denote the j-th term in the above equation for $j\in[5]$, i.e., $T_{x,1}=I,T_{x,2}=|0\rangle\!\langle 0|$, etc. Let $\mathcal{I}$ be the set of all pairs of (x,j), so that $|\mathcal{I}|\leq 5M$. Expanding $\left|J(S^\Phi)\right\rangle^{\otimes q}$ can be written as the linear summation of

$$\bigotimes_{i=1}^{q}(I\otimes|x_i\rangle\!\langle x_i|\otimes T_{x_i,j_i})\,|\Omega\rangle^{\otimes q}$$

of (less than) $(5M)^q$ terms. A similar result holds when expanding $\left\langle J(S^\Phi)\right|^{\otimes q}$, resulting $(5M)^{2q}$ terms total.

Fix one expanded term. It involves at most $2q$ label occurrences. For each label x, the contribution of it is a moment of the form

$$\mathbb{E}\left[|\phi_x\rangle^{\otimes a_x}\langle\phi_x|^{\otimes b_x}\right],\qquad a_x,b_x\leq 2q.$$

If $a_x\neq b_x$, both the Haar moment and the design moment vanish by the phase-invariant property. If $a_x=b_x=s$, then $s\leq 2q$, and the δ-approximate $2q$-design gives

$$\left\|\mathbb{E}_{\psi_x}|\psi_x\rangle\!\langle\psi_x|^{\otimes s}-\mathbb{E}_{\phi_x}|\phi_x\rangle\!\langle\phi_x|^{\otimes s}\right\|_1\leq 2\delta.$$

Since the fixed expanded term contains at most $2q$ relevant labels, a simple hybrid shows that the distance is at most $4q\delta$.

By linearity of expectation, the difference of the two Choi moments is the sum, over all expanded monomials, of the corresponding differences of monomial expectations. Let $\mathsf{J}(A) := (I \otimes A)|\Omega\rangle$. For $\alpha, \beta \in \mathcal{I}^q$, define

$$T_{\alpha,\beta}(\Phi) := \left(\bigotimes_{i=1}^{q} \mathsf{J}(B_{\alpha_i}(\Phi))\right)\left(\bigotimes_{i=1}^{q} \mathsf{J}(B_{\beta_i}(\Phi))\right)^{\dagger}.$$

For each fixed expanded monomial $T_{\alpha,\beta}$, the previous argument shows

$$\left\|\mathbb{E}_{\Phi\leftarrow\mathcal{S}} T_{\alpha,\beta}(\Phi) - \mathbb{E}_{\Psi\leftarrow\mathcal{D}} T_{\alpha,\beta}(\Psi)\right\|_1 \le 4q\delta.$$

There are at most $(5M)^{2q}$ such monomials. Hence, by the triangle inequality for the trace norm,

$$\begin{aligned}
&\left\|\mathbb{E}_{\Phi\leftarrow\mathcal{S}} \left|J(S^\Phi)\right\rangle\!\left\langle J(S^\Phi)\right|^{\otimes q} - \mathbb{E}_{\Psi\leftarrow\mathcal{D}} \left|J(S^\Psi)\right\rangle\!\left\langle J(S^\Psi)\right|^{\otimes q}\right\|_1 \\
&\le \sum_{\alpha,\beta} \left\|\mathbb{E}_{\Phi\leftarrow\mathcal{S}} T_{\alpha,\beta}(\Phi) - \mathbb{E}_{\Psi\leftarrow\mathcal{D}} T_{\alpha,\beta}(\Psi)\right\|_1 \\
&\le (5M)^{2q} \cdot 4q\delta,
\end{aligned}$$

thus the overall bound becomes

$$O\left(4q\delta \cdot (5DM)^{2q}\right) = O\left(q \cdot C^{q^2} \cdot \delta\right)$$

for some $C > 0$, because $D = 2^{O(q)}$ and $M \le 2^{q+1}$. This proves the main upper bound. The final statement is clear because the proof only ever uses independence among the at most $2q$ labels appearing in a fixed expanded term. □

6 On Separating PRUs from PRFSGs

In this section, we consider the length-ℓ quantum-accessible unitarized CHFS oracle $\mathcal{S} = \mathcal{S}_\ell$ for $\ell(|x|) = |x|$ and the **QPSPACE** oracle. We will prove the following theorem, which is the main result of this section.

Theorem 6.1. *There exist adaptively-secure quantum-accessible PRFSGs but there do not exist non-adaptive PRUs whose implementation uses neither ancillary registers nor intermediate measurements, relative to* $(\mathcal{S}, \mathbf{QPSPACE})$.

Formally, the PRUs whose implementation uses neither ancillary registers nor intermediate measurements can be specified by an alternating application of unitaries and oracle queries. The existence of the adaptively-secure quantum-accessible PRFSGs relative to the oracles is proven by Theorem 4.4. What remains is to prove that PRUs without ancilla do not exist in this model.

Lemma 6.2. *Non-adaptive PRUs whose implementation uses neither ancillary registers nor intermediate measurements do not exist with probability 1 relative to the oracle* $(\mathcal{S}, \mathbf{QPSPACE})$.

Proof. We prove the lemma by contradiction. Assume that $\{G_\lambda^{\mathcal{S},\mathbf{QPSPACE}}(\cdot)\}_\lambda$ is a secure $n(\lambda)$-PRU construction relative to $(\mathcal{S}, \mathbf{QPSPACE})$ for $n(\lambda) = \omega(\log\lambda)$. For simplicity, we drop the **QPSPACE** oracle and λ in notations and write $G_k^{\mathcal{S}}$ to denote $G_{|k|}^{\mathcal{S},\mathbf{QPSPACE}}(k)$. The adversary is given oracle access to the oracle $(V, \mathcal{S}, \mathbf{QPSPACE})$ where V is either $G_{k^*}^{\mathcal{S}}$ for some k^* or a Haar random unitary of the same size, and tries to determine which is the case with non-negligible probability.

We write $\mathcal{S} = (S_d)_{d\in\mathbb{N}}$ where $S_d = \sum_{x\in\{0,1\}^d} |x\rangle\langle x| \otimes S_{|\phi_x\rangle}$ to denote the unitary CHFS oracle, where $S_{|\phi_x\rangle}$ denotes the swap oracle defined in Definition 4.1 for some d-qubit Haar random quantum state $|\phi_x\rangle$ and S_d acts on a $(2d+1)$-qubit space.[31]

Let $m = \mathsf{poly}(\lambda)$ be the maximum number of oracle queries to $\mathcal{S}$ that $G^{\mathcal{S}}$ makes. We show that distinguishing $G_k^{\mathcal{S}}$ from a Haar random unitary can be done efficiently based on swap tests. More concretely, we prove that the following algorithm $\mathcal{A}^{V,\mathcal{S},\mathbf{QPSPACE}}$ can guess with non-negligible probability whether V is $G_k^{\mathcal{S}}$ for some random k, (in which case it outputs 1), or a truly Haar random unitary (in which case it outputs 0). For simplicity, we omit the oracle notation and write $\mathcal{A}$ for $\mathcal{A}^{V,\mathcal{S},\mathbf{QPSPACE}}$.

Algorithm 1. $\mathcal{A}$ chooses an n-qubit Haar random state $|\rho\rangle$ and does the following on input oracle access to V.[32]

1. $\mathcal{A}$ executes the purity test $16\lambda^2$ times on $V(|\rho\rangle\langle\rho|)$, which makes $32\lambda^2$ non-adaptive queries to V. If the tests fails at least 4λ times, $\mathcal{A}$ returns 1 and abort. Otherwise, it proceeds to the next step.

2. Let $\tau = 2\log(16m) < n$. For all $i \le \tau$, $\mathcal{A}$ runs $\mathtt{Tom}$ as defined in Theorem 3.19 on the oracles S_i with parameters $\varepsilon = \frac{2}{2^{\tau/2}}, \delta = \frac{1}{2^{2\lambda}}$ and obtains S'_i that approximates S_i. Then it defines a new simulated oracle
$$\tilde{S}_d \coloneqq \begin{cases} I & \text{if } d > \tau, \\ S'_d & \text{otherwise.} \end{cases}$$
We write $\tilde{\mathcal{S}} = (\tilde{S}_d)_{d\in\mathbb{N}}$. Let $r = 1200\lambda$. For each k, we define the following sub-protocol P_k that takes as input a state Ψ over the register $\mathbf{A}_1\mathbf{A}'_1 \dots \mathbf{A}_r\mathbf{A}'_r$:

   P_k**:** Apply $(G_k^{\tilde{\mathcal{S}}} \otimes I)^{\otimes r}(\Psi)$, where each $G_k^{\tilde{\mathcal{S}}}$ acts on $\mathbf{A}_i$. Then, apply the swap test on $\mathbf{A}_i\mathbf{A}'_i$ for each $i \in [r]$. Return 1 if at least $2r/3 = 800\lambda$ tests passes, and return 0 otherwise.

   Note that computing $G_k^{\tilde{\mathcal{S}}}$ does not require any queries to the CHFS reflection oracle $\mathcal{S}$, neither does P_k.

3. $\mathcal{A}$ prepares the following state
$$\Psi \coloneqq \bigotimes_{i\in[r]} \left(|\rho\rangle\langle\rho|_{\mathbf{A}_i} \otimes V(|\rho\rangle\langle\rho|)_{\mathbf{A}'_i}\right).$$
   $\mathcal{A}$ applies Lemma 3.15 on input state Ψ and the family of POVMs induced by $\{P_k\}_{k\in\{0,1\}^\lambda}$, and outputs the same output as the OR tester.

The following claims summarize the main analysis of the algorithm, which will be proven at the end of this section.

Claim 6.3. *Algorithm 1 is a* $\mathbf{BQP}^{V,\mathcal{S},\mathbf{QPSPACE}}$ *algorithm, which makes non-adaptive queries to* V.

Claim 6.4. *If* $V = G_k^{\mathcal{S}}$ *for some* k *and* $\mathrm{Tr}\left(G_k^{\mathcal{S}}(\rho)^2\right) \ge 1 - 1/\lambda$, *then* $\Pr[P_k(\Psi) \to 1] \ge 1 - 2^{-\lambda}$ *holds with probability at least* $1 - \frac{m+\tau}{2^{2\lambda}}$ *over the randomness of the algorithm for sufficiently large* λ.

Claim 6.5. *If* $V \leftarrow \mu_n$, *then* $\Pr[P_k(\Psi) \to 1] \le 2^{-2\lambda}$ *holds for all* k *with probability at least* $1 - 2^{-\lambda}$ *over the randomness of the algorithm for sufficiently large* λ.

The efficiency of the algorithm relative to $\mathcal{S}, \mathbf{QPSPACE}$ is provided by Claim 6.3. Also note that the algorithm breaks the non-adaptive security of PRUs, as the queries to V only occur in the first step and to prepare Ψ which are all non-adaptive queries.

[31] In the proof below, we consider the oracle queries to S_d. The same proof can be extended to the oracle queries to $S_{|\phi_x\rangle}$ for each x, or more general cases. e.g., queries to $|0\rangle\langle 0| \otimes S_{|\phi_{x_0}\rangle} + |1\rangle\langle 1| \otimes S_{|\phi_{x_1}\rangle}$ for any x_1, x_2 of the same length. We focus on the queries to S_d because it is the most complicated.

[32] To be efficient, $\mathcal{A}$ can use s-design for large s instead of Haar random state.

The correctness of the algorithm can be shown by the case analysis. If $V = G_k^{\mathcal{S}}$ for some k and if $\mathrm{Tr}\big(G_k^{\mathcal{S}}(\rho)^2\big) \leq 1 - 1/\lambda$, the first step of $\mathcal{A}$ outputs 1 with probability at least $1 - 2^{-\lambda}$ as shown in Lemma 3.11.

The other case, i.e., $V = G_k^{\mathcal{S}}$ and $\mathrm{Tr}\big(G_k^{\mathcal{S}}(\rho)^2\big) \geq 1 - 1/\lambda$ or V is a true Haar random unitary is dealt with by the quantum OR lemma. In this case, by Claim 6.4 and Claim 6.5, the POVMs $\{P_k\}_{k\in\{0,1\}^\lambda}$ and Ψ satisfy the conditions of the quantum OR lemma (Lemma 3.15) unless with probability $\frac{m+\tau}{2^{2\lambda}} + 2^{-\lambda} \leq 2/2^\lambda$ for large enough λ, $\varepsilon = 1/2^\lambda$ and $\delta = 1/2^{2\lambda}$. Therefore, $\mathcal{A}$ outputs 1 with probability at least $1/8$ if $V = G_k^{\mathcal{S}}$ for some k, but it outputs 1 with probability at most $4/2^\lambda$ if $V \leftarrow \mu_n$, that is, $\mathcal{A}$ breaks the PRU security of $\{G_k^{\mathcal{S}}\}$. This concludes the proof. □

We now prove the claims.

Proof of Claim 6.3. The first step takes polynomial time and $32\lambda^2$ non-adaptive queries to the oracle V for the $16\lambda^2$ purity tests. The second step has time and query complexity (to $\mathcal{S}$) equal to $\tau \times \mathsf{poly}\left(d, \frac{1}{\varepsilon}, \log\frac{1}{\delta}\right) = \mathsf{poly}\left(2^\tau, 2^\tau, \lambda\right) = \mathsf{poly}(\lambda)$. Note that it is clear that P_k can be executed by a quantum polynomial space machine. In the final step, the quantum OR tester can be executed by a **QPSPACE**-aided BQP machine as noted in "Moreover" part of Lemma 3.15 with inputs the descriptions of S'_i for $i \leq \tau$ (prepared by the first step) as P_k can be implemented by a quantum polynomial-space machine. □

Proof of Claim 6.4. We will show that $G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)$ and $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$ are close with high probability, for a Haar random input state $|\rho\rangle$ of size n-qubit. We will write $\rho = |\rho\rangle\langle\rho|$ as a short-hand.

We begin with the following two closeness properties for S_d and $\tilde{S}_d$ from the later steps of the algorithm. First, for small dimensions $d \leq \tau < n$, Theorem 3.19 ensures that

$$\|\tilde{S}_d \otimes I(\rho) - S_d \otimes I(\rho)\|_{tr} = \|S'_d \otimes I(\rho) - S_d \otimes I(\rho)\|_{tr} \leq \frac{\varepsilon}{2} = \frac{1}{2^{\tau/2}}, \tag{29}$$

holds for any quantum state ρ with probability $1 - \delta = 1 - \frac{1}{2^{2\lambda}}$. Thus all tomography outputs are $1/2^{\tau/2}$-close to the target unitaries with overwhelming probability $1 - p_1$, for $p_1 = \tau/2^{2\lambda}$. In the following, we assume it is the case.

Now we discuss the large dimensions $d > \tau$, where we want to show that S_d acts almost as the identity with high probability for a pure Haar quantum state $|\rho\rangle$. Observe that for the reflection $R = I - 2\Pi$ such that $\|\Pi\,|\rho\rangle\,\|^2 = p$,

$$\|\rho - R(\rho)\|_{tr}^2 = \|\,|\rho\rangle\langle\rho| - R\,|\rho\rangle\langle\rho|\,R^\dagger\|_{tr}^2 = 1 - |\,\langle\rho|\,R\,|\rho\rangle\,|^2 = 4p(1-p) \leq 4p.$$

Given this, for $S_d = I - 2\sum_{x\in\{0,1\}^d} |x\rangle\langle x| \otimes |\phi_x-\rangle\langle\phi_x-| \otimes I_{n-2d-1} = I - 2\Pi_d$ for the rank 2^{n-d-1} projector Π_d,

$$\mathbb{E}_\rho\|\tilde{S}_d \otimes I(\rho) - S_d \otimes I(\rho)\|_{tr}^2 = \mathbb{E}_\rho\|I(\rho) - S_d \otimes I(\rho)\|_{tr}^2 \leq 4\mathbb{E}_\rho\|\Pi_d\,|\rho\rangle\,\|^2 \leq \frac{1}{2^{d-1}} \leq \frac{1}{2^\tau}$$

where we use Lemma 3.4. Note that $\|\Pi_d\,|\rho\rangle\,\|^2$ is the success probability of an 1-query algorithm to a Haar random unitary U by letting $\rho = U\,|0\rangle$. Applying Corollary 3.2 with $t = 1/2^{\tau+1}$, we have

$$\Pr_\rho\left[\|\Pi_d\,|\rho\rangle\,\|^2 \geq \frac{1}{2^\tau}\right] \leq \exp\left(-\frac{2^n - 2}{24 \cdot 2^{2\tau+2}}\right) \leq \frac{1}{2^{2\lambda}}$$

for sufficiently large n. In other words, except for the probability $1/2^{2\lambda}$, it holds that

$$\|\tilde{S}_d \otimes I(\rho) - S_d \otimes I(\rho)\|_{tr} \leq \frac{2}{2^{\tau/2}} = \varepsilon \tag{30}$$

To bound the trace distance between $G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)$ and $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$, we use the hybrid argument using the above two observations. Let Φ_j for $0 \leq j \leq m$ be equal to the outcome of G_k on input $|\rho\rangle\langle\rho|$ with the first j

oracle queries are answered using $\tilde{\mathcal{S}}$ and the other $m-j$ queries are answered using $\mathcal{S}$. We have that Φ_0 is the state $G_k^{\mathcal{S}}(\rho)$ and Φ_m is the state $G_k^{\tilde{\mathcal{S}}}(\rho)$.

Let $|\phi_j\rangle$ be the intermediate state right after j-th oracle query when computing $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$. Note that for large d, we apply the identity instead of S_d, thus the intermediate state is independent of S_d. Therefore we can use Eq. (30) for large d with high probability over ϕ_j. For small d, Eq. (29) ensures the upper bound. Thus, by the subadditivity and monotonicity of the trace distance, we have

$$\begin{aligned}\left\|G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)-G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)\right\|_{tr} &\le \sum_{j=0}^{m-1}\|S_{d_j^{(k)}}\otimes I(\phi_j)-\tilde{S}_{d_j^{(k)}}\otimes I(\phi_j)\|_{tr}\\ &\le \sum_{j=0}^{m-1}\max(\frac{\varepsilon}{2},\varepsilon)=\frac{2m}{2^{\tau/2}}=\frac{1}{8},\end{aligned}$$

with probability $1-p_2$ for $p_2=\frac{m}{2^{2\lambda}}$; we again focus on this case.

We finally analyze the success probability of a single swap test between $G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)$ and $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$ succeeds in subroutine P_k. Since

$$\|G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\otimes G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)-G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\otimes G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\|_{tr}=\|G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)-G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\|_{tr}\le 1/8,$$

we have

$$\left|\frac{1+\mathrm{Tr}\left(G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\right)}{2}-\frac{1+\mathrm{Tr}\left(G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)^2\right)}{2}\right|\le\frac{1}{8},$$

and using the fact that $\mathrm{Tr}\left(G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)^2\right)\ge 1-1/\lambda$, we have

$$\frac{1+\mathrm{Tr}\left(G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)G_k^{\mathcal{S}}(|\rho\rangle\langle\rho|)\right)}{2}\ge\frac{7}{8}-\frac{1}{2\lambda}\ge\frac{3}{4}.$$

Therefore, by Chernoff's inequality, the probability that at least $\frac{2r}{3}$ tests succeed among r swap tests is bounded by

$$1-\exp\left(-\frac{3r}{2\cdot 4\cdot 12^2}\right)=1-\exp\left(-\frac{r}{384}\right)\ge 1-2^{-\lambda}.$$

Overall, if $V=G_k^{\mathcal{S}}$ for some k and $\mathrm{Tr}\left(G_k^{\mathcal{S}}(\rho)^2\right)\ge 1-1/\lambda$, then it holds that $\Pr[P_k(\Psi)\to 1]\ge 1-2^{-\lambda}$ with probability at least $1-p_1-p_2=1-\frac{m+\tau}{2^{2\lambda}}$. □

Proof of Claim 6.5. In this case, we can regard $V(|\rho\rangle\langle\rho|)$ as an independent Haar random pure state $|\psi\rangle$. By Lemma 3.10, the expected success probability of the swap test between $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$ and $V(|\rho\rangle\langle\rho|)$ is

$$\mathop{\mathbb{E}}_{V\leftarrow\mu_n}\left[\frac{1+\mathrm{Tr}\left[G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)V(|\rho\rangle\langle\rho|)\right]}{2}\right]=\mathop{\mathbb{E}}_{\psi\leftarrow\sigma_n}\left[\frac{1+\mathrm{Tr}\left[G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)\,|\psi\rangle\langle\psi|\right]}{2}\right]=\frac{1}{2}+\frac{1}{2^{n+1}}$$

where we use Lemma 3.4 in the last equality. Applying Corollary 3.2 for $t=1/13$, we have

$$\Pr\left[\frac{1+\mathrm{Tr}\left[G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)V(|\rho\rangle\langle\rho|)\right]}{2}\ge\frac{7}{12}\right]\le\exp\left(-\frac{2^n-2}{4056}\right)\le\frac{1}{2^{2\lambda}}$$

for sufficiently large n. In other words, with probability at least $1-\frac{1}{2^\lambda}$, the swap test between $G_k^{\tilde{\mathcal{S}}}(|\rho\rangle\langle\rho|)$ and $V(|\rho\rangle\langle\rho|)$ succeeds with probability at most $7/12$ for all k by the union bound. We only focus on such a case below. Chernoff inequality gives that $\Pr[P_k(\Psi)\to 1]$ is at most

$$\exp\left(-\frac{7r/12\cdot(1/12)^2}{(2+1/12)}\right)=\exp\left(-\frac{7r}{3600}\right)\le 2^{-2\lambda},$$

for all k. Therefore, if V is truly Haar random unitary, then it holds that $\mathrm{Tr}\left[P_k \left|\Psi\right\rangle\right] \leq 2^{-2\lambda}$ for all k with probability at least $1 - 2^{-\lambda}$. □

7 On Length Extension of PRSGs

In this section we show that the output size of a pseudorandom state may be relevant, i.e., there exist short-PRSGs but PRSGs in a certain form do not exist.

We will consider a particular type of quantum algorithms, which we call "ancilla-resetting quantum algorithms", where a quantum algorithm A acts on two registers: the output register $\mathbf{A}$, and the ancilla register $\mathbf{B}$ [33], and the output of A is of the following form:

$$A(\cdot) = \mathrm{Tr}_{\mathbf{B}}(\rho_{\mathbf{AB}}), \text{ where } \rho_{\mathbf{AB}} = \psi_{\mathbf{A}} \otimes \left|0\right\rangle\!\left\langle 0\right|_{\mathbf{B}}. \tag{31}$$

Remark 7.1. We focus on algorithms that reset the ancilla to their initial values. More generally, we allow any algorithm that applies only reversible computation to the ancilla, i.e., maps it to a state independent of the oracle. In this case, one can assume without loss of generality that the ancilla are uncomputed back to $\left|0\right\rangle$ at the end of the computation.

We now show the following theorem, which is the main result of this section.

Theorem 7.2. *There exists an isometry oracle $\mathcal{O}$ relative to which (classical-accessible) short-PRFSGs exist, but long-PRSGs with ancilla-resetting generation algorithms making non-adaptive oracle queries and using only constant-size ancilla do not.*

The separating oracle $\mathcal{O}$ consists of two oracles: the classical-accessible isometry CHFS oracle O_ℓ for $\ell(\lambda) = \lfloor 2\log\lambda \rfloor$ and the **QPSPACE** oracle. The existence of short-PRFSGs follows immediately from Theorem 4.4. It remains to break long PRSGs with ancilla-resetting generation algorithm making non-adaptive oracle queries and using only constant-size ancilla.

Proof of Theorem 7.2. By contradiction, assume there exists a PRSG $\mathsf{Gen}(\cdot)$ with an ancilla-resetting generation algorithm relative to $\mathcal{O}$. Let u_k be the length of the ancilla register. We can assume w.l.o.g. that $u_k = u = \mathcal{O}(1)$ is independent of k, by considering $u = \max_k u_k$ and adding ancilla that will not be used for the k such that $u_k < u$. Because the generation algorithm is ancilla-resetting, the ancilla registers are reset to $\left|0\right\rangle$ after the computation. We write $d(\lambda)$ and $\kappa(\lambda)$ to denote the output length and the key length of the PRSG. Since the **QPSPACE** oracle is unitary, we can embed them in the unitaries and write the output state of the algorithm (before tracing out the ancilla) by

$$\widetilde{\rho}^{(k)} = U_k \left(\bigotimes_{i=1}^{t_k} O_{x_i^{(k)}} \right) \left| 0^{m(\lambda)} \right\rangle, \tag{32}$$

where $m(\lambda) = d(\lambda) + u$ is the number of qubits of the whole space where the computations are made, and $x_1^{(k)}, \ldots, x_{t_k}^{(k)}$ are chosen before any oracle query. We omit the superscript (k) when it is clear from the context.

Let us denote by $\widetilde{\rho}^{(k)} = \rho^{(k)} \otimes \left|0\right\rangle\!\left\langle 0\right|^{\otimes u}$ the final state before tracing out the ancilla. We consider the following adversary $\mathcal{A}$, given the polynomial copies of either $\rho = \rho^{(k)}$ for some k (in which case it outputs 1) or Haar random state ρ (in which case it outputs 0). In the following, let $r = C\kappa(\lambda)^2$ for a sufficient large universal constant C and $T \geq 100r^2(2td+1)^3$ where $t = \max_k t_k$.

Algorithm 2. $\mathcal{A}$ does the following on input multiple copies of a state ρ.

1. $\mathcal{A}$ executes the purity test $16T\kappa(\lambda)$ times on ρ. If the test fails at least $4\kappa(\lambda)$ times, $\mathcal{A}$ returns 1 and aborts. Otherwise, it proceeds to the next step.

[33] Wlog, the input register can be part of $\mathbf{A}$ and $\mathbf{B}$.

2. We define the following sub-protocol P_k that takes as input a state $\Psi = ((\rho \otimes |0\rangle\langle 0|^{\otimes u})^{\otimes 2})^{\otimes r}$:

   P_k**:** For each $i \in [r]$, compute $U_k^\dagger \otimes U_k^\dagger(\rho \otimes |0\rangle\langle 0|^{\otimes u} \otimes \rho \otimes |0\rangle\langle 0|^{\otimes u})$ and apply the product test for $\ell(|x_1|), \ldots, \ell(|x_{t_k}|), 1, \ldots, 1$ qubits, where the number of 1 is $s_k = d(\lambda) + u - \sum_{i\in[t_k]} \ell(|x_i|)$. Let $m_k = t_k + s_k$ be the total number of swap tests used in the product test. Return 1 if all tests pass, and return 0 otherwise.

   Then $\mathcal{A}$ runs the quantum OR tester with $\{P_k\}_{k\in\{0,1\}^\kappa}$ on $\Psi = ((\rho \otimes |0\rangle\langle 0|^{\otimes u})^{\otimes 2})^{\otimes r}$, and returns the same output.

We note that the sub-protocol P_k can be implemented in polynomial time.

Claim 7.3. *If $\rho = \rho^{(k)}$ and $\mathrm{Tr}(\rho^2) \geq 1 - 1/T$, then $\Pr[P_k(\Phi)] \geq 4/5$.*

Claim 7.4. *If ρ is a Haar random state, then $\Pr[P_k(\Phi) \to 1] \leq 1/2^{2\kappa(\lambda)}$ for all k with probability at least $1 - 1/2^{\kappa(\lambda)}$.*

The same argument as in Section 6 concludes the proof. Indeed, if $\rho = \rho^{(k)}$ for some k and $\mathrm{Tr}(\rho^2) \leq 1-1/T$, then Lemma 3.11 asserts that the first step outputs 1 with probability $1 - 2^{-\kappa(\lambda)}$.

The other case, i.e., $\rho = \rho^{(k)}$ and $\mathrm{Tr}(\rho^2) \geq 1 - 1/T$ or ρ is a true Haar random state is dealt by the quantum OR lemma. In this case, by Claim 7.3 and Claim 7.4, the POVMs $\{P_k\}_{k\in\{0,1\}^{\kappa(\lambda)}}$ and Ψ satisfy the conditions of the quantum OR lemma (Lemma 3.15) unless with probability $1/2^{\kappa(\lambda)}$. Therefore, $\mathcal{A}$ outputs 1 with probability at least $1/20$ if $\rho = \rho^{(k)}$ for some k, but it outputs 1 with probability at most $4/2^{\kappa(\lambda)} = \mathsf{negl}(\lambda)$ if $\rho \leftarrow \nu_n$, that is, $\mathcal{A}$ breaks the PRSG security of $\mathsf{Gen}(\cdot)$. □

Proof of Claim 7.3. By the purity-test analysis, we can assume that $\mathrm{Tr}(\rho^2) \geq 1 - 1/T$, otherwise Algorithm 2 would have terminated at step 1 with probability at least $1 - 2^{-\kappa(\lambda)}$. Note that $\Pr\left[P_k((\tilde{\rho}_t^{(k)})^{\otimes 2r}) \to 1\right] = 1$ by Lemma 3.12. This implies that P_k outputs 1 on input $\Phi = (\rho^{(k)} \otimes |0\rangle\langle 0|^{\otimes u})^{\otimes 2r}$ with probability at least

$$1 - \frac{1}{T} \geq 4/5.$$ □

Proof of Claim 7.4. Since the CHFS output length is $\ell(n) = \lfloor 2\log n \rfloor$, we therefore have $\ell|x_i| \leq 2\log|x_i| \leq 2c_G \log\lambda$, for some constant c_G for all sufficiently large λ. We thus have

$$m_k = t + s_k \geq \frac{t \cdot 2c_G\log\lambda + s_k}{2c_G\log\lambda} \geq \frac{d(\lambda)+u}{2c_G\log\lambda} = \frac{\omega(\log\lambda)}{2c_G\log\lambda} = \omega(1),$$

where we used the fact that the candidate PRSG has output dimension $d(\lambda) = \omega(\log\lambda)$. Together with the hypothesis that $u(\lambda) = \mathcal{O}(1)$, for all sufficiently large λ and all keys k, this shows that there exists a constant $m_0(u)$ such that $4^u \cdot 2 \cdot (3/4)^{m_0(u)} \leq 0.05$.

By Lemma 3.13 and Lemma 3.14, we have that a single product test (for key k) succeeds with expected probability at most $4^u \cdot 2 \cdot (3/4)^{m_k} \leq 0.05$. By the concentration inequality, we can show that with probability at least $1 - 1/2^{2\kappa(\lambda)}$ over Haar random states, a single product test for k succeeds with probability at most 0.1. Using Chernoff's inequality, we conclude that for each k, $\Pr[P_k(\Phi) \to 1] \leq 1/2^{2\kappa(\lambda)}$. Finally, taking a union bound over all $2^{\kappa(\lambda)}$ keys k, we obtain that the above bound holds for every $k \in \{0,1\}^{\kappa(\lambda)}$ except with probability at most $2^{\kappa(\lambda)} \cdot 2^{-2\kappa(\lambda)} = 2^{-\kappa(\lambda)}$. This proves the claim. □

References

[AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 333–342. ACM Press, June 2011. 5

[Aar04] Scott Aaronson. Limitations of quantum advice and one-way communication. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 320–332. IEEE, 2004. 17

[Aar16] Scott Aaronson. The complexity of quantum states and transformations: from quantum money to black holes. *arXiv preprint arXiv:1607.05256*, 2016. 17

[ABGL25] Prabhanjan Ananth, John Bostanci, Aditya Gulati, and Yao-Ting Lin. Pseudorandomness in the (inverseless) haar random oracle model. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 138–166. Springer, Cham, May 2025. 6

[AE07] Andris Ambainis and Joseph Emerson. Quantum t-designs: t-wise independence in the quantum world. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 129–140, 2007. 29

[AGKL24] Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 226–254. Springer, Cham, May 2024. 7

[AGL24] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common Haar state model: Feasibility results and separations. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 94–125. Springer, Cham, December 2024. 6, 7

[AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 237–265. Springer, Cham, November 2022. 3

[AIK22] Scott Aaronson, DeVon Ingram, and William Kretschmer. The acrobatics of BQP. In *Proceedings of the 37th Computational Complexity Conference*, pages 1–17, 2022. 5

[ALY24] Prabhanjan Ananth, Yao-Ting Lin, and Henry Yuen. Pseudorandom strings from pseudorandom quantum states. In Venkatesan Guruswami, editor, *ITCS 2024*, volume 287, pages 6:1–6:22. LIPIcs, January / February 2024. 3, 4, 6, 7

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Cham, August 2022. 3

[Aub98] Thierry Aubin. *Some nonlinear problems in Riemannian geometry*. Springer Science & Business Media, 1998. 24

[Bar26] Mohammed Barhoush. Separating pseudorandom generators from logarithmic pseudorandom states. In *Advances in Cryptology – ASIACRYPT 2026*, 2026. To appear. 4, 7

[BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997. 4, 21, 45

[BBO+25] Mohammed Barhoush, Amit Behera, Lior Ozer, Louis Salvail, and Or Sattath. Signatures from pseudorandom states via $\perp$-PRFs. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part VIII*, volume 16252 of *LNCS*, pages 320–349. Springer, Singapore, December 2025. 3, 5, 7

[BCN25] John Bostanci, Boyang Chen, and Barak Nehoran. Oracle separation between quantum commitments and quantum one-wayness. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 3–22. Springer, Cham, May 2025. 7

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *ITCS 2023*, volume 251, pages 24:1–24:21. LIPIcs, January 2023. 5, 7

[BFV20] Adam Bouland, Bill Fefferman, and Umesh V. Vazirani. Computational pseudorandomness, the wormhole growth paradox, and constraints on the AdS/CFT duality (abstract). In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 63:1–63:2. LIPIcs, January 2020. 6

[BM24] Samuel Bouaziz--Ermann and Garazi Muguruza. Quantum pseudorandomness cannot be shrunk in a black-box way. *arXiv preprint arXiv:2402.13324*, 2024. 3, 4, 6, 7

[BMM+25] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. A new world in the depths of microcrypt: Separating OWSGs and quantum money from QEFID. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 23–52. Springer, Cham, May 2025. 7, 8, 14, 17

[BNY25] Mohammed Barhoush, Ryo Nishimaki, and Takashi Yamakawa. MicroCrypt assumptions with quantum input sampling and pseudodeterminism: Constructions and separations. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part VIII*, volume 16252 of *LNCS*, pages 516–548. Springer, Singapore, December 2025. 4, 5, 6, 7

[BS20] Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 417–440. Springer, Cham, August 2020. 3, 7

[BT06] Andrej Bogdanov and Luca Trevisan. Average-case complexity. *Foundations and Trends® in Theoretical Computer Science*, 2(1):1–106, 2006. 44

[Bus82] Peter Buser. A note on the isoperimetric constant. In *Annales scientifiques de l'École normale supérieure*, volume 15, pages 213–230, 1982. 24

[CCS25] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: Constructing and separating quantum pseudorandomness. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 108–137. Springer, Cham, May 2025. 6, 7, 10, 14, 17

[CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 215–248. Springer, Cham, August 2024. 3, 4, 5, 6, 7

[CGG+25] Bruno Cavalar, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. On the computational hardness of quantum one-wayness. *Quantum*, 9:1679, 2025. 6

[CM24] Lijie Chen and Ramis Movassagh. Quantum merkle trees. *Quantum*, 8:1380, June 2024. 6

[DLS24] Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail. Fiat-shamir for proofs lacks a proof even in the presence of shared entanglement. *Quantum*, 8:1568, 2024. 6

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986. 3

[GLMY25] Aditya Gulati, Yao-Ting Lin, Tomoyuki Morimae, and Shogo Yamada. Black-box separation between pseudorandom unitaries, pseudorandom isometries, and pseudorandom function-like states. Cryptology ePrint Archive, Paper 2025/1864, 2025. 7

[GMMY26] Eli Goldin, Tomoyuki Morimae, Saachi Mutreja, and Takashi Yamakawa. CountCrypt: Quantum cryptography between QCMA and PP. In *Theory of Cryptography – TCC 2026*, 2026. To appear. 7

[GZ25] Eli Goldin and Mark Zhandry. Translating between the common haar random state model and the unitary model. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 269–300. Springer, Cham, August 2025. 8

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 3

[HKOT23] Jeongwan Haah, Robin Kothari, Ryan O'Donnell, and Ewin Tang. Query-optimal estimation of unitary channels in diamond distance. In *64th FOCS*, pages 363–390. IEEE Computer Society Press, November 2023. 10, 18

[HLM17] Aram W Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1598–1611. SIAM, 2017. 10, 17

[HM10] Aram Wettroth Harrow and Ashley Montanaro. An efficient test for product states with applications to quantum Merlin-Arthur games. In *51st FOCS*, pages 633–642. IEEE Computer Society Press, October 2010. 11, 15

[HY25] Minki Hhan and Shogo Yamada. Pseudorandom function-like states from common haar unitary. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part III*, volume 16270 of *LNCS*, pages 134–165. Springer, Cham, December 2025. 6

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Cham, August 2018. 3, 13

[KN96] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry, volume 2*, volume 2. John Wiley & Sons, 1996. 24

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1589–1602. ACM Press, June 2023. 6, 20

[KQT25] William Kretschmer, Luowen Qian, and Avishay Tal. Quantum-computable one-way functions without one-way functions. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 189–200. ACM Press, June 2025. 5, 6

[Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. 3, 4, 6, 9, 12, 29, 33

[Lic58] André Lichnerowicz. *Géométrie des groupes de transformations*, volume III of *Travaux et Recherches Mathématiques*. Dunod, Paris, 1958. 24

[Lub78] Elihu Lubkin. Entropy of an n-system from its correlation with a k-reservoir. *Journal of Mathematical Physics*, 19(5):1028–1031, 05 1978. 15

[LV24] Romi Levy and Thomas Vidick. Prs length expansion, 2024. 6

[Mec19] Elizabeth S Meckes. *The random matrix theory of the classical compact groups*, volume 218. Cambridge University Press, 2019. 12

[Mel24] Antonio Anna Mele. Introduction to haar measure tools in quantum information: A beginner's tutorial. *Quantum*, 8:1340, 2024. 16

[MH25] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 806–809. ACM Press, June 2025. 7

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Cham, August 2022. 3

[MY24a] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, pages 4–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024. 6

[MY24b] Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part V*, volume 14924 of *LNCS*, pages 359–392. Springer, Cham, August 2024. 6

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991. 3

[NC10] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010. 11

[Qia24] Luowen Qian. Unconditionally secure quantum commitments with preprocessing. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 38–58. Springer, Cham, August 2024. 6

[Rit23] Manuel Ritoré. The isoperimetric profile of compact manifolds. In *Isoperimetric Inequalities in Riemannian Manifolds*, pages 127–155. Springer, 2023. 24

[Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394. ACM Press, May 1990. 3

[Wee06] Hoeteck Wee. Finding pessiland. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 429–442. Springer, Berlin, Heidelberg, March 2006. 4

[Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 80–91. IEEE Computer Society, 1982. 20

[Zha12] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, Berlin, Heidelberg, August 2012. 31

[Zha25] Mark Zhandry. How to model unitary oracles. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 237–268. Springer, Cham, August 2025. 5

A Formal Proof of the Barrier Theorem

A.1 Removing differentiability assumption

In this section, we show how to prove Eq. (16) without assuming the differentiability of $v(r) = \sigma(S_0[r])$. Recall that the Cheeger constant in Eq. (13) and the condition $\sigma(S_0, S_1) \geq \Gamma$ implies that

$$\sigma^+(S_0[r]) \geq h_X \Gamma$$

for $r \in [0, \Delta)$ as in Eq. (15). Let $\varepsilon > 0$. The outer Minkowski measure says that for each r, there exists δ_r such that for any $h \in [0, \delta_r)$ and $r + h \leq \Delta$,

$$\frac{\sigma(S_0[r+h]) - \sigma(S_0[r])}{h} \geq (h_X \Gamma - \varepsilon) \tag{33}$$

where we use $S_0[r][h] = S_0[r+h]$ in our ambient space X. Define the set

$$G_\varepsilon := \{t \in [0, \Delta] : \sigma(S_0[t]) - \sigma(S_0) \geq t \cdot (h_X \Gamma - \varepsilon)\}.$$

We will prove that $r \in G_\varepsilon$ for all $r < \Delta$. Assuming this, we have

$$\sigma(S_0[r]) - \sigma(S_0) \geq r \cdot (h_X \Gamma - \varepsilon) \tag{34}$$

thus choosing $\varepsilon \to 0$ proves Eq. (16).

Let $a := \sup G_\varepsilon \cap [0, r]$, which is well-defined because 0 is included in G_ε. Let us assume that $a \in G_\varepsilon$, which will be proven later. If $a < r$, for a sufficiently small h such that $a + h \leq r$, we have

$$\begin{aligned}
\sigma(S_0[a+h]) - \sigma(S_0[0]) &= (\sigma(S_0[a+h]) - \sigma(S_0[a])) + (\sigma(S_0[a]) - \sigma(S_0[0])) \\
&\geq h \cdot (h_X \Gamma - \varepsilon) + a \cdot (h_X \Gamma - \varepsilon) \\
&= (h + a) \cdot (h_X \Gamma - \varepsilon)
\end{aligned}$$

where, to prove the inequality, we use Eq. (33) in the first term and the fact that $a \in G_\varepsilon$ in the second term. This says that $a + h \in G_\varepsilon \cap [0, r]$, which contradicts to $a = \sup G_\varepsilon \cap [0, r]$. Therefore $a = r \in G_\varepsilon \cap [0, r]$, proving Eq. (34).

It remains to prove that $a \in G_\varepsilon \cap [0, r]$. By the definition of a, there exists an increasing sequence $(a_n)_{n \in \mathbb{N}}$ that converges to a. Since $\sigma(S_0[r])$ is an increasing function, we have

$$\sigma(S_0[a]) - \sigma(S_0) \geq \sigma(S_0[a_n]) - \sigma(S_0) \geq a_n \cdot (h_X \Gamma - \varepsilon)$$

where the last inequality holds because $a_n \in G_\varepsilon$. Taking $n \to \infty$ gives

$$\sigma(S_0[a]) - \sigma(S_0) \geq a \cdot (h_X \Gamma - \varepsilon),$$

proving $a \in G_\varepsilon$. This concludes the proof.

B Average-case Hard Problem

This section presents an average-case hard samplable promise problem in **QCMA** $\cap$ **coQCMA** relative to the log-length unitary CHFS oracle with a sketch of the analysis. The non-existence of QOWF is clear due to the main theorem. All results in this section is in the relativized world.

We first formally define the samplable problem distributions and related quantum complexity, adapting [BT06].

Definition B.1. *We say that a probability ensemble $\{D_n\}_{n\in\mathbb{N}}$ is (QPT) samplable if there is a quantum polynomial time sampler S such that $\Pr[S(1^n) \to x] = D_n(x)$ for all n and $x \in \{0,1\}^*$. We denote by* **sampQCMA** $\cap$ **sampcoQCMA** *the class of distributional promise problems consisting of decisional promise problems in* **QCMA** $\cap$ **coQCMA** *coupled with samplable probability ensembles whose outputs satisfy the promise with an overwhelming probability.*

We say that a samplable decisional problem $(\mathcal{L}, \{D_n\})$ is quantumly easy-on-average if there exists a BQP algorithm A, such that

$$\Pr_{x\leftarrow D_n}\left[\Pr_A[A(x) \to \mathcal{L}(x)] \le 2/3\right] = \mathsf{negl}(n).$$

Otherwise, we say that $(\mathcal{L}, \{D_n\})$ is hard-on-average. Our goal is to exhibit a samplable problem $(\mathcal{L}, \{D_n\})$ in **sampQCMA** $\cap$ **sampcoQCMA** that is hard-on-average. Note that the existence of quantum-computable one-way functions implies the existence of a quantumly hard-on-average problem in **sampQCMA** by the Goldreich-Levin theorem.[34] The converse of this implication does not hold in our world, where quantum-computable one-way functions do not exist, yet there exists a quantumly hard-on-average problem in **sampQCMA**, suggesting that our world is a quantum-computable version of Pessiland.

We consider the unitary CHFS oracle model with $\ell(\lambda) = \lceil c \log \lambda \rceil$ for sufficiently large $c \gg 1$. In other words, our CHFS oracle maps $|x, 0, 0\rangle$ to $|x, 1, \phi_x\rangle$ for $\ell(|x|)$-qubit state ϕ_x. We have the following properties for the CHFS oracles.

- The output states are almost orthogonal. More precisely, with probability 1 over the choice of the CHFS oracle, $|\langle\phi_x|\phi_{x'}\rangle|^2 \le 0.01$ holds for all $(x, x') \in \{0,1\}^\lambda \times \{0,1\}^\lambda$ for all sufficiently large λ. This is because of Lemma 3.4 and Corollary 3.2.[35]
- The output states admit polynomial-length classical descriptions. More precisely, there are two efficient procedures $\mathsf{Description}$ that maps a (polynomial copies of) $\ell(n)$-qubit quantum state to a $\mathsf{poly}(n)$-length classical string and $\mathsf{Construct}$ that maps a $\mathsf{poly}(n)$-length classical string to an $\ell(n)$-qubit quantum state such that:
$$\|\rho - \mathsf{Construct}(\mathsf{Description}(\rho))\|_1 \le 2^{-|\rho|}.$$
Furthermore, $\mathsf{Construct}$ is deterministic.[36]

Given these observations, we define the distribution ensembles that sample from D^{yes} with probability $1/2$ and from D^{no} with probability $1/2$. Here we use $\le$ to denote the lexicographical order.

(D^{yes}) We pick a random pair $x \le z$, and outputs $(\mathsf{Description}(\phi_x), z)$.

(D^{no}) We pick a random pair $x > z$, and outputs $(\mathsf{Description}(\phi_x), z)$.

Define

$$\Pi_n := \{(y,z) : \exists! x \in \{0,1\}^n \text{ such that } \|\mathsf{Construct}(y) - \phi_x\|_1 \le 0.01\}.$$

For promised inputs $(y, z) \in \Pi$, let $x(y)$ denote this unique string. We define the language $\mathcal{L}$ so that $\mathcal{L}(y, z) = 1$ only if $x(y) \le z$. Equivalently,

$$\Pi_{yes} = \{(y,z) \in \Pi : x(y) \le z\}, \qquad \Pi_{no} = \{(y,z) \in \Pi : x(y) > z\}.$$

By definition, this language is in **sampQCMA** using such an x as the witness. It is also clear that $\mathbf{Supp}(D^{yes}) \subset \Pi_{yes}$ and $\mathbf{Supp}(D^{no}) \subset \Pi_{no}$ except with negligible probability due to the tomography errors.

To analyze, we recall that $|\phi_x\rangle$ is almost orthogonal to $|\phi_x\rangle$ for $x' \ne x$, which implies that for each y, there is at most a single x such that $\|\mathsf{Construct}(y) - \phi_x\|_1 \le 0.01$. This shows that x can be a witness for **coQCMA** as well, showing that the language is in **sampcoQCMA**.

The proof of the quantum average-case hardness can be directly proven using the argument analogous to the proof of the PRFSG security in the main body and the BBBV lower bound for unstructured search [BBBV97].

[34] Similarly, the one-way permutation implies a hard-on-average samplable problem in **sampQCMA** $\cap$ **sampcoQCMA**.
[35] This requires $c \gg 1$ when using the union bound.
[36] The description can be an approximation of all amplitudes.